nature portfolio

Peer Review File

Laser writing in glass for dense, fast and efficient archival data storage

Corresponding Author: Dr James Clegg

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Referee #1

(Remarks to the Author)
In this study, the authors present an optical storage technology that utilizes femtosecond laser direct writing in silica glass. The technology demonstrates remarkable advancements in practical data storage applications, featuring exceptionally long data lifetimes, automated large-scale data writing, and efficient reading and encoding methods. While femtosecond laser-based optical storage in glass materials has been extensively explored in the literatures, this work distinguishes itself through its system integration and exhibits good performance in back-end data retrieval.
After a thorough review of the manuscript, I have compiled several comments and questions that I believe the authors should address to strengthen the manuscript further:

1: At line 182, The authors state that when reading voxels, "lower NA values do not sufficiently resolve signals from adjacent voxels, while higher NA values introduce more significant spherical aberration with depth." This issue is analogous to what occurs in femtosecond laser writing, where aberrations arise in the direction of laser focusing as the writing depth changes. Specifically, in a large depth range such as 2 mm, this would inevitably result in voxel shape distortions at varying depths. Beyond energy compensation, is there a need for additional aberration correction? Should the authors provide further analysis on the processed voxels within the 2 mm depth range to address this concern?

2: The data densities presented in the manuscript appear to be the best theoretical values. When the writing density is increased (e.g., 300 layers, large area preparation), the overall transmittance of the sample decreases, which could compromise the accuracy of reading. Are the data presented in Figure 4 derived under monolayer structure conditions at a specific depth, or were they tested in samples with higher data densities? The reading difficulty for such samples would likely differ significantly.

Furthermore, when voxels are written at large scales, essentially forming a grating structure, does this induce light diffraction, thereby increasing reading errors? It is mentioned in Section 6.4 that diffraction occurs at specific angles—does this effect change with large-scale writing?

3: This study involves multiple voxel-writing principles, all of which have been thoroughly reported in previous works. For example, the other researchers employ a single-pulse preparation method in glasses, crystals, and highly flexible polymers. To better highlight the novelty of this work, the authors should elaborate on the distinct advantages of their writing method in comparison to other existing techniques.

4. In line 322, The authors claim that "choosing a write laser wavelength from near-infrared to visible and ultraviolet can further enhance writing efficiency and storage density." Unlike photolithography, femtosecond laser direct writing primarily relies on nonlinear multiphoton absorption. Thus, the authors should clarify whether the use of ultraviolet wavelengths can indeed improve writing efficiency and storage density. Additionally, as storage capacity increases, higher requirements for data reading are inevitable. What is the limiting capacity for resolvable read data in the context of this work?

5. The citation information in the References section appears incomplete. In the current manuscript PDF, references lack

essential publication details such as page numbers or article identifiers. Additionally, some cited references are outdated and do not adequately represent recent advancements in the field. For instance, “Such modifications have been observed in various types of transparent materials 26,27 at lower cost than fused silica or quartz glass.”—a claim here that is too narrow to substantiate the broader research context. In reality, femtosecond laser-induced refractive index (RI) modification has been extensively utilized for waveguide fabrication across diverse dielectric materials, representing a well-established technique.

To better contextualize this work, I recommend incorporating recent review articles that comprehensively summarize the state-of-the-art progress in this field. While I have not listed specific reviews here, a number of high-quality and relevant articles have been published and would significantly enhance the background discussion for readers and researchers.

6. The authors employed fused silica and borosilicate glass as optical data storage media. It is important to note that the fs-laser modification characteristics in dielectric materials are highly dependent on their intrinsic properties. Significant differences exist between crystalline and amorphous materials in their response to fs-laser processing. Even within glass systems, the material composition plays a critical role in determining the modification behavior.

The results presented in this work also demonstrate that the fs-laser processing parameters optimized for silica cannot be directly extrapolated to other glass types, as evidenced by the distinct performance differences between fused silica and borosilicate glass. This material-dependent behavior should be explicitly addressed and slightly discussed in the manuscript to provide proper context for the observed results.

7. Given this paper's submission to Nature, I propose including a demonstration of basic optical data storage and retrieval using real-image examples (e.g., in the supplementary materials) to enhance cross-disciplinary appeal.

In conclusion, while the novelty of the voxel-writing processing methods using femtosecond laser direct writing in glass is not particularly groundbreaking, the work demonstrates systemic results in large-scale writing, designing systems for reading, and optimization using machine learning. Regardless of whether it is suitable for publication in Nature, the research is highly significant and contributes valuable insights to the field.

Referee #2

(Remarks to the Author)
- Summary of the key results

Authors present an extended development, with large incremental advancements, in many technical fronts, in the field of volume data storage in dielectrics with laser writing.

The basic science is not new, but technical advancements are very significant and cover diverse aspects that are all needed in order to produce an industrial product, from the well known phenomena of permanent index change writing and reading inside transparent media such as glass and crystals, as well as birefringent voxel writing.

Key results span various engineering strategies to write faster in the glass, like:
- Writing with a continuous scan rather than stop-and-go.
- On-the-fly writing of 2 to 4 spots for both index-change and birefringent-marking, with smart strategies in the case of double-pulse writing of birefringent voxels (authors should change "pseudo-pulse" to "double-pulse", since the word "pseudo" adds zero information).
- Efficient pulse splitting and beam modulation and laser scanning schemes (although a severe lack of precise optical setup beam handling elements diagrams plagues the manuscript, which does not allow to see the optical setups in detail).
- Technical solutions to improve fast reading of distorted lattices by means of 2D marking codes which a reader can use to compensate the mentioned geometrical distortions.
- Optimisation of pulse thermal fingerprints overlap by means of simulations of 3D thermal heat diffusion, well known in the field.
- Writing baseline pre-optimisation strategies based on measuring focal volume emission across line scans, and presetting compensation or "flatening" power-vs-position curves, with closed-loop control.
- and machine learning decoding protocols for fast and precise reading.

- Originality and significance: if not novel, please include reference

The fundamental knowledge of 3D laser-writing permanent index change and utilising it to produce a scheme for writing and reading, is know for 3 decades.
In fact, the full area of research on "3D femtosecond-pulse laser writing" actually was initially motivated for data storage. So, it has been a motivational workhorse for generations of researchers, now.
Nowadays, it is a fully commercial technology.

Without going into details, key bibliography is:

- E. N. Glezer, M. Milosavljevic, L. Huang, R. J. Finlay, T.-H. Her, J. P. Callan, and E. Mazur, "Three-dimensional optical storage inside transparent materials," Opt. Lett. 21, 2023-2025 (1996)
- Watanabe, M., Juodkazis, S., Sun, H. B., Matsuo, S., Misawa, H., Miwa, M., & Kaneko, R. (1999). Transmission and photoluminescence images of three-dimensional memory in vitreous silica. Applied physics letters, 74(26), 3957-3959.
- Zijlstra, P., Chon, J. & Gu, M. Five-dimensional optical recording mediated by surface plasmons in gold nanorods. Nature 459, 410–413 (2009). https://doi.org/10.1038/nature08053
- Zhang, J., Gecevičius, M., Beresna, M., & Kazansky, P. G. (2014). Seemingly unlimited lifetime data storage in nanostructured glass. Physical review letters, 112(3), 033901.

- Data & methodology: validity of approach, quality of data, quality of presentation.

All methodologies seem valid and are reasonably documented.
Not much technical details are given on optical parts for achieving efficient beam scan (homogeneous). Most optical step-us sketchs (Supplm. Figs. 11, 20, 24 and 33, for example, have many "black-boxes" of key elements).

- Authors make an appropriate use of statistics and treatment of uncertainties

- Conclusions seem valid and reliable.

- There is little room for suggesting improvements or requiring more data, since authors have been very extensive in details, apart from the mentioned exclusion of optical systems specifications and schemes.

- Bibliographic references seem appropriate, although I am not an expert in the specific field of laser data storage.

- The abstract should be written much more technical, and authors should be much more honest explaining how old the idea of permanent 3D data storage and optical reading, in transparent dielectrics, is. Authors should be more specific on their technical advancements, since as it is written, they seem to appropriate well know ideas in the field, like permanent data storage and reading for thousands of years, Mbit/s writing speeds, and the use of galvo scans and optimization strategies for laser pulse delivery in synchrony with laser-matter interaction and influence pulse dose per unit of time and space.

Referee #3

(Remarks to the Author)
The manuscript is a highly detailed description of a nearly completely novel optical data storage system. Besides the novelty of the complete system, the work contains a number novelties that could be useful to others in a variety of contexts. The authors have revealed details on every aspect of the data storage process and, together, indicate a viable data storage method from the scientific/technology perspective. While several of the results have been revealed in earlier publications, they have been put in the context of the entire system and thus remain important here. The results are highly detailed and credible and indicate a complete proof-of-principle of a new data archiving technology and a blueprint for further development. They reflect the broad expertise of the large team of authors. Of particular importance are the results of data integrity, error correction and important specifications such as capacity and speed. The ability to deliver data at the MB/second rate puts the technology in the company of existing optical data storage methods. The small nonlinear response and consequently relatively high writing threshold are like to limit the speed. In fact the main novelty is not silica or glass, it is the system for data storage in a three-dimensional monolithic responsive material, of which glass is one of the less responsive (but long-lived) materials.

Other, more responsive materials could close the speed gap compared to existing magnetic technology and lower the cost of equipment. Indeed, the novelty here is not necessarily the media, but the details of the data storage system. For example, the authors have demonstrated a high read rate using machine learning image processing methods far different that the signal channel processing of incumbent storage media…a major breakthrough that will simplify data recovery by eliminating the need for nanoscale-precision servos.

Another major innovation is high degree of multiplexing in having several possible symbols to be written in a single voxel. This important novelty is worthy of a short section in the main manuscript, rather than buried only in the supplementary material. In fact, in reading the main manuscript it was not immediately apparent what mechanism for multiple symbols were and how more bits could be written is a single voxel. They were briefly mentioned when listing the capacity and speed.
The word emission emerges on page 21 without context. It became apparent that the emission was used for control. The authors should describe the source of the emission. This reviewer assumes that the emission is nonlinear in the incident intensity…is it? If so, how does this affect its function?
While this reviewer is not expert in error correction, the detailed description of the correction and data integrity and quality was reassuring.

In summary, this manuscript is a tour-de-force exposition of a novel data archiving system, and perhaps more important, presents a number of novelties that could be important in other optical data storage contexts. Publication is strongly recommended; workers in the field are likely to be keenly interested.

Referee #4

(Remarks to the Author)
This paper describes the optical recording techniques used in the Silica digital data recording system. The system uses femtosecond laser pulses to record information in glass in two ways: phase voxels (in borosilicate) and birefringent voxels (in silica). Written data is then read using optical microscopy, with machine learning used to decode the written information back into digital form. The paper studies multiple aspects of this approach, including volumetric density of data, durability, and energy efficiency, showing the tradeoffs in these metrics as optical parameters change.

This work is highly original, leveraging existing work in laser optics and machine learning to build a full system that can record data that will remain durable for millennia. While the recording mechanism is complex, the reading mechanism is more straightforward, making it more likely that the written data can be read at a (much) later date. The paper doesn't describe the full stack, but that information is available elsewhere (e.g., ACM Transactions on Storage 21(1), Feb 2025). Of particular interest is the analysis of the tradeoffs between recording parameters and density and quality of the recording; the paper does a good job explaining how different parameters such as spacing affect the recorded information. The paper also makes (to me) a surprising discovery: data at the "lowest" level is of lower quality not because of the write process, but because of the intervening glass on read. This finding will be useful for future systems that can leverage this knowledge to do double-sided recording.

The experiments in the paper are thorough, though much information is omitted, presumably due to space requirements. Conclusions are well-supported by the experiments. The experiments were run multiple times, including the read-back of written data to ensure that it was recorded properly.

The organization of the paper is somewhat disjointed, but that may be due to my unfamiliarity with Nature's format requirements. Rather than including all of the material (essentially) linearly, information is split among many sections. For example, descriptions of the write process can be found in Sections 2 and 6; it might be helpful to merge the information into a single section. If that was done because of article requirements in Nature, however, it would be acceptable to leave it that way. Certainly, the explanation of experimental techniques in Section 6 is valuable, so it should remain available to readers. However, Section 6 also includes information about the system itself and its metrics; for example, Section 6.2 contains details about the system construction. Similarly, Section 6.7 contains details about the coding techniques used to encode data into symbols; this information is crucial to understanding the system itself, not merely the experiments that were performed. Again, it's possible that the standard format for Nature relegates this detail to a separate section, but it would be very helpful if it were collocated with the earlier section.

Similarly, Appendix A provides a good deal of useful information but, as with Section 6, would be better integrated into the body of the paper rather than relegated to the end. If this is standard practice for Nature, this advice may be ignored, but it would make the paper easier to read and follow if the material were presented in a more cohesive manner.

Section A.2.2 raises an interesting issue: basic writer sweeps involve multiple adjustments of the recording parameters. Does this variation depend on the properties of the glass itself, and, if so, what additional tuning might be required for different glass manufacturers? While this information doesn't directly impact the experiments reported in the paper, it is of paramount importance to a commercial system based on this technology.

The references are sufficient for the paper.

Version 1:

Reviewer comments:

Referee #1

(Remarks to the Author)
The author has made satisfactory revisions and has addressed all of my comments. I recommend acceptance as is.

Referee #3

(Remarks to the Author)
The authors have successfully addressed the concerns listed in the initial review. The statistical analysis is appropriate for the functional prototype reported in the manuscript.

We thank all the referees for their thorough and insightful comments. Their suggestions have helped improve the clarity and presentation of our key findings, and we hope the updated manuscript will resonate with Nature's broad readership.

Below, we provide detailed point-by-point responses to all referees' comments. The referees requested changes the text to clarify a number of technical points, which we have been able to do while editing down the main section to just under the required 3500 words.

Referee 1

Item 1: At line 182, The authors state that when reading voxels, "lower NA values do not sufficiently resolve signals from adjacent voxels, while higher NA values introduce more significant spherical aberration with depth." This issue is analogous to what occurs in femtosecond laser writing, where aberrations arise in the direction of laser focusing as the writing depth changes. Specifically, in a large depth range such as 2 mm, this would inevitably result in voxel shape distortions at varying depths. Beyond energy compensation, is there a need for additional aberration correction? Should the authors provide further analysis on the processed voxels within the 2 mm depth range to address this concern?

Yes, this is a good insight. We use a spherical aberration correction collar on the writing and reading objectives. The collar is actuated by a software-controlled motorised mechanism. This was designed in-house. We verified that this system works well by measuring code rate as a function of depth in the glass. These results are shown in Figure 3e, taken together with the observation that when reading the glass upside down we see the quality as a trend reverse, which demonstrates that the variation we see in voxel quality with depth is due to a reading effect.

We have clarified the main text and methods to say that any time we translate our glass in z, with write and read systems, we adjust the spherical aberration collar of the objective.

Item 2: The data densities presented in the manuscript appear to be the best theoretical values. When the writing density is increased (e.g., 300 layers, large area preparation), the overall transmittance of the sample decreases, which could compromise the accuracy of reading. Are the data presented in Figure 4 derived under monolayer structure conditions at a specific depth, or were they tested in samples with higher data densities? The reading difficulty for such samples would likely differ significantly.

Furthermore, when voxels are written at large scales, essentially forming a grating structure, does this induce light diffraction, thereby increasing reading errors? It is mentioned in Section 6.4 that diffraction occurs at specific angles—does this effect change with large-scale writing?

This is an important question. The densities quoted are not theoretical best values and are not derived from monolayer measurements. For every measurement shown in the main paper (and many more besides), we wrote data throughout the full depth of the glass, read data from a large area preparation across all depths, and fully recovered that data using our decoder.

Item 3: This study involves multiple voxel-writing principles, all of which have been thoroughly reported in previous works. For example, the other researchers employ a single-pulse preparation method in glasses, crystals, and highly flexible polymers. To better highlight the novelty of this work, the authors should elaborate on the distinct advantages of their writing method in comparison to other existing techniques.

We have made changes to try and emphasize the novelty of our work more clearly.

Our data writing methodology addresses two critical challenges that have received limited attention: write throughput and voxel transparency. By transparency, we refer to the ability to inscribe voxels with sufficient structural modification to be reliably detected by the read subsystem, while preserving adequate optical transmittance through a 2 mm glass depth.

To the best of our knowledge, no prior studies have reported the usable data capacity we achieve across a 2 mm-thick glass platter - 2.02 TB for phase voxels and 4.84 TB for birefringent voxels - while operating at the fundamental limit of laser repetition rate (1 pulse for phase and 2 pulses for birefringence). By leveraging a 10 MHz repetition rate and multibit encoding, we attain the highest reported write throughput in both regimes: 18.4 Mbit/s for phase voxels and 25.6 Mbit/s for birefringent voxels.

To underscore the development and implementation of these two write regimes, we have revised the wording in Section 2 as follows:

> "To overcome these limitations, we **developed** two advanced regimes: pseudo-single pulse writing of birefringent voxels and single pulse writing of phase voxels. These approaches maximize energy efficiency and **enable write throughput to reach the laser repetition rate**. The resulting weak modifications exhibit low scattering and minimal crosstalk, enabling diffraction-limited 3D data writing and readout from more than 300 layers (Sec. 3a and b), **improving capacity significantly beyond reported approaches[15,24,26]**.

While the two write regimes employed - birefringent voxels and phase voxels - are not themselves novel, our contribution lies in demonstrating a reliable, end-to-end data storage pipeline encompassing encoding, writing, reading, and decoding, all supported by robust error correction to ensure data integrity. To the best of our knowledge, such a comprehensive demonstration has not been presented in the literature and is essential for the broader deployment of this technology.

Item 4: In line 322, The authors claim that "choosing a write laser wavelength from near-infrared to visible and ultraviolet can further enhance writing efficiency and storage density." Unlike photolithography, femtosecond laser direct writing primarily relies on nonlinear multiphoton absorption. Thus, the authors should clarify whether the use of ultraviolet wavelengths can indeed improve writing efficiency and storage density.

Additionally, as storage capacity increases, higher requirements for data reading are inevitable. What is the limiting capacity for resolvable read data in the context of this work?

Thank you for the question. We have chosen to remove the claim about increasing storage density using different writing wavelengths from the manuscript, as it warrants further investigation beyond the scope of the present study.

To elaborate, we operate within the multiphoton absorption regime. Here the photon energy remains below the bandgap of the material. We initially hypothesised that reducing the wavelength would result in a narrower beam waist, thereby producing smaller voxels and potentially increasing data density. Furthermore, we anticipated that the energy required for voxel formation might decrease, as achieving the same threshold intensity would demand less energy due to the tighter beam focus.

However, we recognise that the relationship between wavelength and write efficiency or density is more nuanced than beam waist considerations alone. Several additional wavelength-dependent factors influence voxel formation, including the number of photons required to bridge the bandgap, the degree of voxel elongation driven by non-linear refractive index changes, and variations in avalanche photoionization.

Considering these complexities, in response to the second part of the Referee's comment, we have demonstrated successful readout at an *xy* voxel pitch of 500 nm × 700 nm using a numerical aperture (NA) of 0.6 on the read subsystem, as shown in Figure 4b. We anticipate that further reductions in voxel pitch could be supported by increasing the NA of the read subsystem and using a shorter wavelength light source for reading.

Item 5: The citation information in the References section appears incomplete. In the current manuscript PDF, references lack essential publication details such as page numbers or article identifiers. Additionally, some cited references are outdated and do not adequately represent recent advancements in the field. For instance, "Such modifications have been observed in various types of transparent materials[26,27] at lower cost than fused silica or quartz glass."—a claim here that is too narrow to substantiate the broader research context. In reality, femtosecond laser-induced refractive index (RI) modification has been extensively utilized for waveguide fabrication across diverse dielectric materials, representing a well-established technique.

To better contextualize this work, I recommend incorporating recent review articles that comprehensively summarize the state-of-the-art progress in this field. While I have not listed specific reviews here, a number of high-quality and relevant articles have been published and would significantly enhance the background discussion for readers and researchers.

We agree that further contextualisation of our study with respect to more recent literature would improve the background discussion. As such, we have additionally cited the following articles in Section 2:

- Lei, Y. et al. "*Ultrafast Laser Writing in Different Types of Silica Glass*". Laser Photonics Reviews **17**, 2200978. doi:10.1002/lpor.202200978 (2023).

- Yao, H., Pugliese, D., Lancry, M. & Dai, Y. "*Ultrafast Laser Direct Writing Nanogratings and their Engineering in Transparent Materials"*. Laser & Photonics Reviews **18**, 2300891. doi:10.1002/lpor.202300891 (2024).
- Wang, H. et al. "*100-Layer Error-Free 5D Optical Data Storage by Ultrafast Laser Nanostructuring in Glass"*. Laser & Photonics Reviews **16**, 2100563. doi:10.1002/lpor.202100563 (2022).

The citations for the line mentioned by the referee have been updated to include two more recent publications, which capture the study of various transparent materials for optical data storage based on isotropic refractive index changes (i.e., phase voxels):

> "Such modifications have been observed in various types of transparent materials[8,28-30]."

Finally, we have also reviewed the References section and updated missing page numbers and DOI article identifiers.

Item 6: The authors employed fused silica and borosilicate glass as optical data storage media. It is important to note that the fs-laser modification characteristics in dielectric materials are highly dependent on their intrinsic properties. Significant differences exist between crystalline and amorphous materials in their response to fs-laser processing. Even within glass systems, the material composition plays a critical role in determining the modification behavior.

The results presented in this work also demonstrate that the fs-laser processing parameters optimized for silica cannot be directly extrapolated to other glass types, as evidenced by the distinct performance differences between fused silica and borosilicate glass. This material-dependent behavior should be explicitly addressed and slightly discussed in the manuscript to provide proper context for the observed results.

We decided to exclude crystalline materials from our investigations, due to their undesirable long-range order effects - such as cleavage and dislocations - which tend to enlarge the modified volume and constrain achievable density. We chose to focus on borosilicate glasses for phase voxel formation, primarily due to the cost-efficiency of large-scale production via float processes compared to high-purity fused silica, while retaining the necessary optical and chemical properties. This consideration is crucial for the commercialisation and broader deployment of this technology.

Although a comprehensive characterisation of various glass types lies beyond the scope of this paper, we fully acknowledge that the composition and atomic structure of the glass significantly influence the voxel formation process. Accordingly, this aspect remains a key focus of our ongoing research efforts.

To address the referee's comment, we have added the following two sentences to the end of Section 4, as follows:

“**The optimal operating parameters for Silica are expected to vary with glass composition[8] since the intrinsic material properties influence voxel formation behaviour. This work should be repeated in order to evaluate alternative glasses.”**

Item 7: Given this paper's submission to Nature, I propose including a demonstration of basic optical data storage and retrieval using real-image examples (e.g., in the supplementary materials) to enhance cross-disciplinary appeal.

We are not exactly sure what is being requested here. We had two possible interpretations of the question:

i. It is possible that the reviewer is asking to see example image data, e.g. of a cat, being stored and retrieved from the glass. If this is what is being asked then we emphasise that this is a digital storage technology with error correction (i.e. not like a microfilm) and so any input, whether it be image, audio, text, etc, will be converted to a stream of bits and that stream of bits will be encoded, stored, read, decoded and recovered. Our decoder only indicates success if all of the original bits are recovered (after error correction has been applied).
ii. It is also possible that the reviewer is asking to see more raw readout images of voxels from our system. We have already shown cropped raw images in figures 3a and 3b, and several more in the extended data figures and believe that adding more raw images would not help the discussion.

If we have misinterpreted the question, then we are happy to receive further feedback.

Referee 2

Item 8: The fundamental knowledge of 3D laser-writing permanent index change and utilising it to produce a scheme for writing and reading, is know for 3 decades.
In fact, the full area of research on "3D femtosecond-pulse laser writing" actually was initially motivated for data storage. So, it has been a motivational workhorse for generations of researchers, now. Nowadays, it is a fully commercial technology.

The referee rightly notes that the foundational principles of 3D femtosecond-pulse laser writing have been extensively explored over the past three decades. Early investigations were primarily driven by the promise of high-density optical data storage - a vision that has significantly influenced the trajectory of the field.

While the underlying physics are now well established, this study presents a comprehensive, end-to-end pipeline that spans data encoding, writing, reading, and decoding. It directly addresses several persistent challenges that have hindered broader deployment. These include substantial improvements in write throughput to commercially viable levels (18.4 Mbit/s for phase voxels and 25.6 Mbit/s for birefringent voxels), refined symbol placement strategies for multi-level encoding, and the integration of robust error-correction mechanisms - specifically LDPC and network coding - to ensure data integrity. To the best of our knowledge, these critical contributions have not been comprehensively addressed in prior literature (see Supp. Table 1), yet they are essential for real-world commercialisation.

In summary, while the historical context is acknowledged and appreciated, the manuscript's core contribution lies in its effort to move beyond foundational physics and toward practical implementation - particularly in the context of ultra-long-term data preservation.

We have addressed similar comments raised by Referee 1 (Items 3 and 5) and have revised the manuscript to more clearly highlight the contributions in relation to prior work in this area.

Item 9: Not much technical details are given on optical parts for achieving efficient beam scan (homogeneous). Most optical step-us sketchs (Supplm. Figs. 11, 20, 24 and 33, for example, have many "black-boxes" of key elements).

We thank the referee for their comment. Full technical details of the optical setups employed in this study are provided in the Methods (Section 6.2) and illustrated in Figure 2a, without the use of black-box representations. To support educational clarity, we included supplementary setup diagrams to demonstrate that certain optical components may be substituted with alternatives. For instance, the acousto-optic deflector (AOD) could be replaced with a polarization grating, a polarization beam splitter, or a spatial light modulator (SLM) for birefringent voxel writing.

However, we acknowledge that the use of generalized black-box representations may introduce ambiguity. Considering editorial constraints on the Supplementary Information (SI), we have opted to remove all additional optical setup diagrams. Specifically, Figures 11, 20, 31, and 33 have been excluded from the revised SI.

Item 10: The abstract should be written much more technical, and authors should be much more honest explaining how old the idea of permanent 3D data storage and optical reading, in transparent dielectrics, is. Authors should be more specific on their technical advancements, since as it is written, they seem to appropriate well know ideas in the field, like permanent data storage and reading for thousands of years, Mbit/s writing speeds, and the use of galvo scans and optimization strategies for laser pulse delivery in synchrony with laser-matter interaction and influence pulse dose per unit of time and space.

We are very keen to emphasise that our advancements in any single metric are not the main contribution of our work, but rather it is the fact that we have been able to simultaneously demonstrate performance across the range of metrics that matter for a real storage system. We make this point in the summary paragraph and in the introduction to the paper. If further emphasis is needed then we are happy to consider it.

Referee 3

Item 11: Another major innovation is high degree of multiplexing in having several possible symbols to be written in a single voxel. This important novelty is worthy of a short section in the main manuscript, rather than buried only in the supplementary material. In fact, in reading the main manuscript it was not immediately apparent what mechanism for multiple symbols were and how more bits could be written is a single voxel. They were briefly mentioned when listing the capacity and speed.

We thank the referee for highlighting the significance of the ability to encode multiple symbol states within a single voxel. For phase voxels we modulate the amplitude of the laser beam to specific values that results in any given voxel having one of N different modification strengths. This means each voxel can store up to log_2(N) bits. For birefringent voxels we modulate the polarization of the writing laser so that each voxel has one of M different azimuth values. This means each voxel can store up to log_2(M) bits.

To increase clarity in the manuscript, we have added two sentences in Section 1 when describing Fig. 1d, as follows:

> "**Bits are then grouped into symbols (Fig. 1c). Each symbol corresponds to one voxel and stores more than one bit. The laser energy or polarization is modulated as the laser beam is moved relative to the glass sample, or platter** (Fig. 1d)."

We have also amended the following innovation claim in Section 1:

> "We demonstrate high quality voxels, **each storing more than one bit**, using only single pulse per phase voxel and a pseudo single-pulse regime for birefringent voxels."

Item 12: The word emission emerges on page 21 without context. It became apparent that the emission was used for control. The authors should describe the source of the emission. This reviewer assumes that the emission is nonlinear in the incident intensity…is it? If so, how does this affect its function?

The use of photo-induced emission in the context of closed-loop feedback control is introduced in Section 2 and elaborated in Section 6.2.3. However, recognizing that this initial mention may be easily overlooked, we have revised the wording in Section 4.1 (page 21) to clarify the implementation:

> "For each voxel pitch experiment, we optimize across the photo-induced emission **used for closed-loop feedback control (see Section 6.3.1)**, the number of symbols per voxel, and the energy modulation for each symbol."

Additionally, we have added a comment to the methods section under emissions-based control to highlight the source of this signal:

> "The write objective is used to collect the emissions profile in a back-reflection configuration. **This emission arises from the generation of a dense and hot plasma during voxel formation[27].**"

The referee also raises a valuable point regarding the relationship between incident pulse energy and photo-induced emission intensity. We have characterized this relationship in borosilicate glass (Schott Borofloat 33), where a near-linear response is observed within the typical writing energy range (14-20 nJ). Glass modification becomes detectable above an energy threshold of 13.7 ± 0.5 nJ for a voxel pitch of (0.5, 0.7) μm in the (x, y) plane.

At higher pulse energies (>25 nJ) and finer voxel pitches, however, deviations from linearity occur due to thermal accumulation effects. Residual heat from adjacent voxel formation reduces emission intensity relative to linear predictions. As such, a deviation from linearity can be the first indication of voxel crosstalk that we wish to avoid.

Item 13: The organization of the paper is somewhat disjointed, but that may be due to my unfamiliarity with Nature's format requirements. Rather than including all of the material (essentially) linearly, information is split among many sections. For example, descriptions of the write process can be found in Sections 2 and 6; it might be helpful to merge the information into a single section. If that was done because of article requirements in Nature, however, it would be acceptable to leave it that way. Certainly, the explanation of experimental techniques in Section 6 is valuable, so it should remain available to readers. However, Section 6 also includes information about the system itself and its metrics; for example, Section 6.2 contains details about the system construction. Similarly, Section 6.7 contains details about the coding techniques used to encode data into symbols; this information is crucial to understanding the system itself, not merely the experiments that were performed.

We have organised our work in accordance with the requirements for Nature, but we thank the reviewer for their helpful comments. Section 6, for example, is the methods section that contains information that is necessary for reproduction of the results we present, but the details do not affect the flow of the main arguments of the paper.

Item 14: Similarly, Appendix A provides a good deal of useful information but, as with Section 6, would be better integrated into the body of the paper rather than relegated to the end. If this is standard practice for Nature, this advice may be ignored, but it would make the paper easier to read and follow if the material were presented in a more cohesive manner.

In response to editorial and referee feedback, we have now split Appendix A into seven Extended Data figures for clarity. These figures are referenced in the main text at appropriate points to aid cohesion but are not essential for understanding the paper. Appendix A now contains only the comparison of storage technologies and supplemental methods.

Referee 4

Item 15: Section A.2.2 raises an interesting issue: basic writer sweeps involve multiple adjustments of the recording parameters. Does this variation depend on the properties of the glass itself, and, if so, what additional tuning might be required for different glass manufacturers? While this information doesn't directly impact the experiments reported in the paper, it is of paramount importance to a commercial system based on this technology.

We thank the referee for their insightful question regarding recording parameters across different glass compositions and acknowledge its relevance to the development of commercial systems. While a detailed analysis of performance variations in different glasses lies beyond the scope of this paper, we are actively investigating this area as part of our ongoing research.

We anticipate that the optimal recording parameters - such as emission setpoint, voxel pitch, focusing conditions, number of symbols per voxel, and energy modulation per symbol - are inherently dependent on the material properties of the glass. For instance, variations in refractive index necessitate adjustments to the spherical aberration correction collar to maintain voxel fidelity throughout the full writing depth (here demonstrated over 2 mm). Similarly, dopants used to tailor the optical absorption properties of the glass can affect the emission setpoint. Thermal properties also play a critical role, particularly in determining voxel separation, as they influence the rate at which heat dissipates post-voxel formation.

Despite the breadth of this parameter space, the workflow outlined in Section 4 offers a versatile framework that can be adapted to a wide range of glass compositions. By systematically optimizing voxel pitch, emission setpoints, the number of symbols per voxel, and energy modulation of each symbol, the methodology accounts for variations in material properties.

We note that this issue has already been addressed in our response to Referee 1, Item 6. Accordingly, we hope this comment to have been sufficiently resolved and believe that no further revisions are required to the manuscript.