Responses to the first review of the Climate DT manuscript

We address the comments by the three reviewers below. We make reference to the parts of the revised manuscript where the corresponding changes can be found.
The reviewers will note that an additional author has been introduced. The additional author was one of the main developers of the AQUA package and has contributed to the monitoring and analysis of the climate simulations.

Reviewer #1

This paper presents an impressive undertaking to address the significant gaps in the ability to provide information from climate models that is user-relevant and accessible. The extent of the technical challenges that have been addressed by this project to fill these needs are incredible. I see this as a framework that could potentially be expanded to S2S predictions as well.
We thank the reviewer for the kind words about the relevance of this manuscript.
The comment about the expansion of this technology to climate prediction is very relevant and will be shared with the programme managers.

There are a few minor suggestions and that I have around making this paper more accessible to a broader audience.
1. There are a lot of acronyms, project names, and technical jargon throughout this paper. Any effort to reduce this would make the paper more accessible.
We have created a table that collects a description of all the manuscript acronyms and another one with the project names. The tables can be found in section 9.

2. It comes across as if this system can do everything in the climate information space. As with anything that is trying to accommodate many users and meet a wide range of needs, I expect there are some limitations and challenges to doing this. Any effort to discuss this would be helpful.
The system cannot solve all the challenges identified by the community in the climate information space. It is clear that it needs to collaborate with the many actors that can facilitate the delivery of robust climate information. Besides, the Destination Earth Service Platform, which is not described in the manuscript because it is beyond the manuscript remit, but is mentioned in it, offers access to both Climate DT data and applications that use them. This data can be combined, if so desired, with other sources of climate information in the Platform. We have modified sections 7 and 8 to describe the limitations of the current system when it comes to providing decision-oriented climate information and the benefits of working together with these other actors.

3. I realize this paper and journal is focused around geoscience model development. I also think the user-focused approach and capabilities is such an important part of this work that it could be expanded a little more here. It seems like a complicated system. How does a user

get training to setup and use the system for their specific needs and get involved in co-production efforts?

The Climate DT system is run by the partnership implementing it. It cannot be run by external users. However, users can influence the way it develops and, as explained in the manuscript, access Climate DT data soon after they are produced. Guidance is available to access all the data included in the DestinE portfolio as documented in the “Code and data availability” section of the manuscript.

Key users were involved in the development of each Climate DT impact sector application described in the manuscript. The developers of each use case were in contact with their respective key users to understand their needs and, as far as possible, take them into account during the system development. This collaboration was very helpful in obtaining feedback about the adequacy of the design. Supporting an increased uptake of Climate DT data, from both applications and climate models, is one of the objectives for the next phase of DestinE.

Besides, examples of how to use the Climate DT are already available on the DestinE Service Platform. These examples are continuously expanded.

We have added more details along these lines in the revised sections 6, 7, and 8.

Reviewer #2

This is potentially a very important paper, detailing the many developments behind the creation of the first European Climate Adaptation Digital Twins, part of the DestinE programme, as well as the overall progress of the DestinE programme against the initial aims and ambitions.
As many in our community, I have been accompanying the DestinE journey, and looking forward to learning about its methods, challenges, achievements, as well as lessons for the future.
While reading the paper, understanding it to be a paper that purports to lay out the technical foundations and going on to current achievements, I continued to expect to learn in good detail about the journey, the challenges, the breakthroughs, the concrete outcomes so far. All of these elements are there, albeit in far less detail than I had expected from a technical paper.
At this time, I find myself wondering whether this is a vision paper, but we have already had a few at the inception of the programme, or a high-level overview paper, aiming to attract new users to the programme, which I think belongs elsewhere, finally whether more could be done with the current manuscript to provide a detailed technical description of what the DestinE Digital Twins, and their enabling technologies are.
I do believe this to be an important paper in principle, so I am going to provide some concrete suggestions to help steer it towards its main objectives.
We thank the reviewer for his constructive comments and for considering the manuscript to be an important contribution. His comments are very useful and will surely help to produce a more readable and informative manuscript.
This is neither a vision (all the elements described already exist and are in use) nor an overview paper (the attempts to attract users are being handled by other means, like the Destination Earth Service Platform managed by ESA). It aims to be an overall description of the Climate DT system. It describes the principles followed, the concept, the essential elements, their links and dependencies, its links to other programme components, and the near-future plans. We thought long and hard about the most adequate journal for this material given the complexity of the system and the intrinsic multidisciplinary character. We concluded that GMD gathers an important part of the readership that we target.
As for the level of detail, there is a limit to what can be provided in the manuscript to produce a readable piece. As the manuscript tries to illustrate, there are many elements (technical, scientific, and social) contributing to the Climate DT system. A number of manuscripts describing key components of the system in more detail have either been submitted or are in the last stages of preparation.

I find the the most important sections are:
SECTION 3: workflows, and the ability to deploy several models and data analyses in a unified way over multiple platforms
SECTION 4: seamless/homogeneous data access
SECTION 5: production of first set of projections and storylines, including their quality assessment
SECTION 6: the impacts applications
In the interest of readability, and of effectively communicating the extremely impressive and important achievements in the sections above, I suggest a substantial reduction to the

introduction, motivation etc., which contains material already available in other papers. I also suggest shortening Section 7, which contains many principles, ambitions etc.. albeit too little in terms of concrete future plans, so not as informative as the previous sections. Can more focussed and more tangible ideas be introduced about future plans, based on the lessons learned so far?
I think that it would strengthen the paper to provide more evidence in Sections 3-6, particularly in terms of why some key decisions were taken, e.g. from the point of view of performance, portability, longevity, sustainability of the solutions.
We have tried to reduce section 1. However, we found making a substantial reduction hard for two reasons: a) not every reader will be familiar with the previous literature about the digital twin concept and b) although there is literature about requirements to produce climate information, the requirements of equitability, timeliness and co-production guide the Climate DT development and, to our knowledge, are not formulated in a similar way in the existing literature.
Section 2 sets the scene. Although it has been shortened too, the operationalisation concept and the Climate DT links to the Destination Earth (DestinE henceforth) infrastructure, including the definition of the Climate DT boundaries in this infrastructure, are needed to contextualise the sections the reviewer singles out as most relevant.
As for section 7, it has been substantially shortened and focused on the existing Climate DT capabilities and lessons learned, as requested.
Finally, sections 3 to 6 have also been considerably rewritten.

I also think that the paper could be improved in a few areas.

1. Concise, relevant and effective communication.

The paper contains many repetitions, particular in regards to its high-level objectives. For instance, the word "equitable" is repeated six times, within sentences that are quite similar to each other, yet the paper provides very little to demonstrate that this goal has been achieved. This principle is a crucially important part of the original aims, but are we at the stage where any progress can be measured? The word "credible", which follows, used 5 times, has more of a chance of being supported by evidence from the programme's achievements so far, and we read about quality assurance, but very little detail is provided, which is a pity. Next is the word salient, 4 times, but there is not much evidence in the paper that the information provided so far is indeed salient. Such value judgements may be applied in the future, once the community has truly adopted the DestinE products, but it seems too early to insist on those at this time.
We agree with the reviewer that it is too early to assess whether these principles have been satisfied. At the same time, we find that it is important for an operational system like the Climate DT to be guided by clear principles inspired by the experience acquired in interactions with communities involved in climate adaptation.
The text does not make reference to the principles the reviewer refers to as Climate DT achievements. Instead, it mentions them as 1) requirements identified by previous literature (for instance, outside the abstract, "credible" was mentioned twice as a previously identified requirement and only once, near the end, as a Climate DT ambition) and 2) principles to guide the system development (for instance, "equitable" was always used as an ambition to satisfy user requirements for climate information rather than a Climate DT achievement).

The end-to-end workflow, the climate model diagnostics, the validation of the application output, the automatic data quality control, and the system traceability are all aspects introduced to contribute to the credibility of the climate information.
We have revised the text to reduce repetitions of those principles and to make clear that they guide both current and future Climate DT developments, although it will take time to assess whether they are fully achieved.
With respect to quality assurance, section 5.1 describes the real-time scientific quality assessment, which has been expanded in the revised manuscript with more illustrations. This is complemented by the description of the data checks and the workflow provenance in sections 3 and 4.

The "summary and conclusions" contains very long bullet points, covering nearly two pages, and repeats much of what had already been said in the sections that come in the main body of the paper. The information presented is, once again, rather programmatic, instead of focussing on the essence of what has been achieved in the project so far.
We have worked hard on this comment. Section 8 has been substantially shortened following the reviewer's recommendation. The text now briefly lists the main achievements and mentions short-term plans inspired by the most recent lessons learned.

Some of the key figures are too small, too low resolution, and use too small fonts. Examples: Figs 1,2, 6.
Figures 1, 2, 3, 4, and 6 have been redrawn, some of them being completely different now. The layout has been changed to make the inserted text more easily readable.
It is worth noting that the low resolution of the figures might be due to the automatic conversion to pdf performed by the submission system. The original figures are all of high quality and follow the journal's specifications.

I suggest:
- stating the high-level aims only once, at the very start of the paper, especially since we have already had papers about the DestinE vision in the past
- summarising progress, based on evidence, in this DestinE phase, against those initial aims
- shortening the summary and conclusions to 1/2 - 2/3 of a page
- make some of the key figures larger and more readable: use them to illustrate key concepts and technologies as soon as possible
We have addressed all these recommendations, hopefully correctly interpreting the reviewer's advice.
As explained before, we agree that there is already literature referring to digital twins in climate so we will reduce the introduction. As this literature does not describe the guiding principles in detail, we have retained this part.

2. Reporting on groundbreaking technical progress

The multiple technological choices made during the development and demonstration phases of the programme are introduced several times over, albeit initially in a qualitative, unspecific way. Later in the paper we start to learn the names of some of the key technologies and what they can do, albeit not the first time they are mentioned, which is rather confusing. The one exception is the discussion of HEALPix, which is all in one place, and quite complete.

We have tried to address this point during the revision, consolidating the description of the choices made. As an example, the key technologies have been listed in figure 2.

I also find that a few important things are left undefined, e.g. "bias adjustment" in FIg 1 and "bias correction algorithm" in Fig 2. There is brief mention of some examples around lines 491-495, but which ones are routinely used in the workflow?

We hope the revised version has managed to address any lack of definition of important aspects. Figure 2 includes a list of the technologies used in the system.

Concerning the references to bias, for bias adjustment we have implemented the method described by Lange (2019, https://doi.org/10.5194/gmd-12-3055-2019), which is a classical quantile mapping approach adapted to be compatible with the streaming concept through the use of the t-digest algorithm for estimating probability density functions. We have used t-digest also to estimate cumulative density functions in the one-pass algorithm. All the details are available in Grayson et al. (2025, https://doi.org/10.5194/gmd-18-5873-2025), which has been referred to in all the relevant parts of the manuscript.

Bias adjustment (BA) and bias correction were used as interchangeable terms to denote the statistical post-processing applied to climate model output to reduce systematic differences relative to a reference observational dataset. To avoid confusion, we have consistently used the term "bias adjustment", which is the term recommended in the last IPCC report, throughout the revised manuscript.

BA in Figure 1 refers to the component of the workflow where model output can be statistically adjusted before being passed to downstream applications, whereas a BA algorithm denotes the specific statistical method and its software implementation used within a task.

The brief mention of "model biases" around line 491 is unrelated to the BA implementation itself and instead refers to the AQUA scientific quality assessment. This has been made clearer in the revised manuscript.

Regarding the question about which approaches are routinely used: the HydroLand application can already be executed with BA. The application has been run with data from the 2025 production cycle with BA activated. BA is a relatively new component in the workflow and will undergo further development as the Climate DT implementation progresses and more climate simulations are performed.

Moreover, there are some clear areas of groundbreaking progress that deserve more prominence:

The aim of the paper is to give an overview of the Climate DT system and its technical choices so far. It would be impossible to cover every aspect of the system in detail. Several publications covering the individual components are in preparation or have already been submitted. This is the case of the OPA, AQUA, the quality of the Climate DT simulations, and the description of the storyline simulations.

- the climate models, which are now more portable, faster, more energy efficient, and, from other evidence I have seen, producing better results, in many areas, than typical CMIP models. Please discuss more detail on these advances than you currently have.

We have increased the material about the climate model validation (now in figures 4 and 6), although we have not yet included a comparison with CMIP simulations. This assessment is still ongoing work and we prefer to have more solid material before claiming that results are better in some aspects. We are aware that the results from related projects like

PRIMAVERA, EERIE, and nextGEMS show an improvement of the model behaviour with resolution and are working to check that the same applies to the Climate DT climate simulations.

- the unified workflow manager: being able to run several models at once through a single workflow manager is obviously attractive. However, what are the detailed tradeoffs in terms of speed, portability, resilience, convenience to operators and end-users? Should such a tool be adopted by the entire community?

It is important to note that the Climate DT's workflow achievements and corresponding benefits are twofold. First, we use only one orchestration component, the Autosubmit workflow manager, to operate the system from start to finish. Second, we employ a single workflow software project to encapsulate all the logic that Autosubmit requires to manage the workflow. The convergence of models and applications was necessary to implement the Climate DT and benefit from the streaming capability.

Regarding the workflow manager choice, as the Climate DT system has been established in a very short period of time, it would be too early to draw lessons for the entire climate modelling community at this stage. Autosubmit is a versatile tool that proved to be suitable for many contexts, from research to operations, from shorter to longer scales. It was critical to operate the system seamlessly in more than one HPC. Autosubmit runs fully in a virtual machine, orchestrating the work in MareNostrum5 and LUMI from a central point. This is clearly convenient for the system operators. But it is hard to say now if either this one, or an alternative workflow manager, could cover all the needs in the community. BSC is taking steps to compare and favour interoperability between Autosubmit and other workflow managers.

The system portability is discussed in section 3. Using one workflow software, containers for the data consumers, and one orchestration tool reduces the number of dependencies, facilitating portability. Section 8 makes reference to modularity and portability, and the plans to extend it to more data consumers. These aspects will be improved in future production cycles.

The text could have probably been more detailed regarding resilience. Autosubmit is a robust workflow manager. It has demonstrated its ability to run operational setups on a set of non-dedicated computing platforms, such as the EuroHPC ones. Having only one workflow manager reduces points of failure and helps addressing common issues. Section 3, where the unified data handling, system configuration, and workflow control are introduced, describes now the system resilience and the benefits of this strategy for the operators.

- part of the above could be showing an example of the output of one of the quality assurance diagnostic/metrics tools, while more detailed scientific analyses should be left for specialised papers.

Figure 4 now shows examples of the typical evaluation results.

- the new data translation and analysis tools, workflows, etc.

We have updated the data flow section and linked it to the revised workflow description.

- the discovery environment, starting with the data catalogues, sample notebooks/scripts, data analysis platforms, and including the AI "chatbot".

The manuscript focuses on the Climate DT system and does not describe the service components that are handled by DestinE's Data Lake and Service Platform. These

components give external users access to the data and examples of how to analyse them. We have tried to make these boundaries clearer in the manuscript, as well as the fact that efforts to develop more examples for data analysis and to support user uptake will grow over time.
As for the AI chatbot, it is described in two publications. One was included in the manuscript and another one (https://www.nature.com/articles/s44168-025-00300-y) has been added in the revised version.

- the platforms for running end-user models: what are they, what do they enable?
We understand the reviewer refers here to the DestinE Service Platform. As mentioned above, this is not the focus of the paper. The Climate DT simulations cannot be run by external users for a number of reasons like the computational cost and access policies to the supercomputing platforms. Should the comment refer to the impact applications embedded in the Climate DT, they run on the same computing platform as the rest of the system elements, in all cases using containerised solutions to facilitate portability. The Climate DT capabilities are evolving, and efforts are planned in the near future to increase the usability of the Climate DT data, in particular as far as its access and the translation of climate data into tailored information for specific sectors is concerned.

I suggest:
- providing full details of each technological solution the first time it is mentioned
- consider building a concise table with the names of the key technologies, what they are built on, what they deliver,
- explaining why these technologies have been chosen over traditional ones: what are the tradeoffs?
- altering Fig2 and Fig4 to contain the names of the key enabler technologies
We think that we have addressed all these recommendations in the revised manuscript. We tried to describe the technologies integrated in the Climate DT, and added a list of them to figure 2, justified their use (with their pros and cons). However, a detailed description of each and every element of the Climate DT system cannot be included in a single manuscript. Many of these details will be included in separate, dedicated submissions.

3. Users, uptake, co-design etc.

I suggest that you inform on:
- How many users have signed up to the services so far?
We appreciate the interest in the impact of the DestinE infrastructure, but this is a question outside the scope of this manuscript and the knowledge of the co-authors. DestinE's users are managed by the DestinE Service Platform, which is a separate component of DestinE's infrastructure.

- How much resource have they used so far, both in terms of CPU and storage?
As in the previous comment, this information is not available to us. DestinE's users can access the Climate DT data via the Service Platform, which is connected to the Data Lake (Climate DT's long-term repository) and is gradually increasing its user base.
If the question refers to the resources used by the Climate DT operation, in the period 2023-2025 operators and developers have used 5% (up to 10%) of the GPU-based (CPU-based)

partitions of the LUMI and MareNostrum5 HPCs. As for the data generated, the data bridges linked to those HPC platforms host at this stage around 8 PB of data, although this number keeps growing.

- What is the capacity of the system(s), now and in the future: how many users can it accommodate, envisaging a range of uses with different requirements?
The manuscript focuses on the system development and operation. External users cannot run the Climate DT system, which is managed by the partnership implementing DestinE. The results are transferred after being produced to the Data Lake and made accessible to external users via the Service Platform. Describing the Platform policy, objectives, and strategy is beyond the scope of this paper. We have described the boundaries of the manuscript in the revised version.

4. Future outlooks

I found a lack of concrete guidance on the next steps. While I think that convincing new users to come to the platforms and use/co-design the DestinE data, I had expected to learn about the technologies developed to support such new interactions, that is, what can prospective users expect to find in terms of resources, tools, longevity of the data, stability and longevity of the platforms themselves?
As mentioned above, the paper describes the Climate DT system, not the entire DestinE ecosystem, i.e., including the DestinE Data Lake and Service Platform. The Platform is the gateway for accessing DestinE data and provides several services and resources to access data and perform data analysis. The Climate DT does not control the resources, tools and guidance made available to Platform users. The longevity of the platforms and Climate DT data transferred depends on the DestinE governance. We have tried to make this clearer in the revised manuscript.

I suggest that you point people to some of the key resources and on how to apply for them, as well as explaining what specific type of support users can expect in the future. Finally, it would be important to draw a roadmap for users to be in the position of effectively influencing/designing future experiments.
Access to DestinE data can be obtained by applying via the DestinE Platform by all the categories of users mentioned at https://platform.destine.eu/general/updated-ec-terms-conditions-extension-of-upgraded-access-policy/. This will be made clearer in the revised manuscript.
Until now, stakeholder engagement has taken place through intensive discussions with key users in order to gather their feedback. As implementing this feedback often requires a substantial amount of time, it cannot easily scale. Developers take advantage of DestinE's events like the regular User Exchange meetings (https://destine.ecmwf.int/5th-destination-earth-user-exchange/), where many members of the Climate DT team actively participate. Based on this experience and development, the circle of stakeholders and feedback is expanded. Further feedback is submitted via the DestinE website and portal and reaches the developers through ECMWF's officers. However, the extent to which user suggestions can be taken into account depends on the complexity of the level of development required, the social relevance of the issue to be resolved, the number of users requesting it, and the resources available.

Defining the roadmap for users to further influence DestinE is beyond the scope of this manuscript. It depends on the negotiations between the delegated entities. However, the development team is making efforts to facilitate the users' role. For instance, a key feature of the ClimateDT is the implementation of what-if scenarios motivated by user requirements: what if a certain adaptation measure (e.g., levee construction, green infrastructures in cities) are implemented? From a technical point of view, what-if scenarios are modified configurations of the ClimateDT workflow, either the end-to-end one or the workflow limited to the data consumers (i.e., accessing data from the data lake). While modifying climate model configuration is a challenging task, especially in an operational setting, what the Climate DT has achieved is the possibility to implement what-if options for the integrated impact applications. Examples of these possibilities include the consideration of different turbine types for renewable energy application according to key user feedback or forest cover changes within the hydrology application.

Specific points

Line 111: do also mention what carrying such responsibilities on behalf of the community means for individual career progression, and how these advances can break that cycle.
Following another reviewer's comment, we have removed that part of the sentence. Besides, expanding on this specific point would make the introduction even longer when we have been advised to focus on making it shorter and more focused.

Line 117: can you be more specific when you say what could be done to complement CMIP and CORDEX? For instance, do you mean generation of short bursts of large ensembles at HR using emulators empowered by ML?
This is one option currently under discussion that is being explored in a companion DestinE contract. Other options that emerged from discussions between CMIP and Copernicus where the Climate DT was invited is a more timely production of both climate forcings and scenarios, as explained in https://www.nature.com/articles/d41586-025-02642-3. As far as we know, Copernicus and DestinE are teaming up with the relevant institutions to make this option a reality, so that the Climate DT projections fulfill the requirements formulated by policy-making institutions like the European Environment Agency. The Climate DT team is now working on a comparison between the Climate DT simulations and the equivalent ones from CMIP and CORDEX. We cannot include these examples in the revised manuscript because all of them are either still under development or under discussion, but emphasised the complementarity.

Lines 159-164: you talk about trustworthy, but as far as we know many of the model codes are not open-source. From what is stated, the workflows are fully accessible, so does this mean that users may ask to replicate experiments and/or test them for robustness? More generally, what are the key ingredients of this trustworthiness?
Trustworthiness refers to a system that is both reproducible, traceable, and produces quality-assured information. Reproducibility consists in establishing operational, version-controlled workflows with components that can reproduce the simulations (should the same computing platforms for the climate models be available), similarly to what is done in operational weather predictions. As the reviewer points out, not all the software is openly available, but it is well identified (as requested by the journal). Provenance information is centrally stored in

RO-Crate files to support traceability. Quality assurance is supported by the scientific evaluation of the climate simulations and application output, as well as by the data checks and growing documentation of the system. Yet, a user might not test the system for robustness because external users do not have access to it, as mentioned in the responses to previous comments.
We will include this aspect in section 8 of the revised manuscript.

Line 185: there are two concepts: timely and routine production: what are their exact definitions in this (currently semi-operational, going towards operational) context?
Routine refers to the aim to produce regular climate projections with updated Climate DT components. The frequency is at least once per year, depending on the availability of computing resources. Timely refers to a production that obeys a predefined schedule, where system infrastructure developments, climate models, scenarios, and data consumers, among other components, are integrated to allow for the routine production by dedicated operators. We have tried to make these concepts clearer in the revised manuscript.

Line 187: Unprecedented: can you quantify and put in context?
The Climate DT, with several climate models performing global, high-resolution simulations simultaneously and offering hourly output for the atmosphere and land and daily output for the ocean and sea ice, produces several hundreds of TBs per week. The sentence refers to the centralised management of these data volumes. The order of magnitude of the data volume production rate can be found in section 4.

Lines 189-: you start to take about data streaming, efficient data handling, what does it mean in practice? What is the baseline, and what are the requirements for DestinE to be declared a success?
We consider that drawing conclusions about the global relevance and level of success of the data streaming is preliminary given the stage of implementation and the so far reduced exposure of the Climate DT.
The baseline would be an approach where a selection of climate variables are output by the climate models and transferred to a long-term storage to make it available to data consumers after the climate simulations have been completed. Given the data volumes produced in an operational setting, this selection would have to be relatively limited or, if large, data be thinned otherwise (e.g., by reducing resolution). Data streaming has been chosen as an alternative approach to consume data directly from the climate model production. This approach takes into account both the large data volumes produced and the constraints that apply to the storage available in supercomputing platforms. Data consumers are encouraged to process the climate data as it is produced (and quality checked) so that they have access to the best estimate possible of the model state vector for a limited period of time. Of course, this requires a high level of abstraction and efficiency, which is supported by the GSV, the data notifiers, the OPA, and the common workflow. Data can be thinned (e.g., interpolated, averaged, compressed) after being consumed before transferring it for long-term preservation and exploitation in a different infrastructure (the DestinE Data Lake and Service Platform).
We hope that the description of the streaming to the embedded data consumers is now clearer in the revised manuscript.

Line 195: “Design OF adaptation strategies”

Thanks, this has been changed in the revised version.

Line 205: DT infrastructure is its flexibility: what does this mean in practice? Is the following text just a definition and/or vocational, or have you implemented those things? It is not clear whether this paragraph is still introductory, or reporting on what has been accomplished. There are some undefined concepts, such as "fast testing of climate adaptation options (line 211). What does it mean concretely?
The characteristics mentioned in the paragraph have been implemented, each one being in a different stage of maturity. The text summarised some of the current capabilities. It stated that the system is flexible meaning that it is portable, modular, adaptable, and scalable, as described in some of the responses above.
We understand that it could be premature to advance these characteristics in this section without offering the corresponding context. Hence, we have simplified the paragraph making clear that these are capabilities described in detail in the subsequent manuscript sections.

Lines 212 and 237: these are examples of instances in which tools are mentioned, albeit not named nor defined, so the sentences read as rather vague and uninformative.
We have modified these sentences to make the text more concrete.

Line 264: is "Autosubmit" the tool you were mentioning before? It is the only one? Where does it stand internationally? Should it be adopted by others? Can it?
Yes, it is Autosubmit and is the only workflow manager used by the Climate DT. We believe it is both preliminary and beyond the scope of this paper to draw such conclusions. What we can share is that Autosubmit is used in European projects like ESiWACE3 and has been adopted by several European institutions (for instance, in the ICON community LMU, KIT, and MPI-Jenna) to orchestrate their workflows. For instance, it is the workflow manager of choice in the digital twin of the ocean implemented in EDITO.

Line 294: you introduce the idea of speed of access. What is the speed required to enable all the DestinE objectives, and what has been achieved so far, both in local node modality and for distributed access?
This sentence has been removed. The minimum speed required is determined by the ability of all the data consumers to use the data before it is transferred and erased to make space for additional model output.

Line 299: Particular features are: particular features of what? Also, what is "lazily loaded"?
It refers to the Climate DT data management. This text was clearly confusing. We have completely modified the sentence and removed these terms.

Line 313: over-precision: do you mean over-resolution? Resolution does not necessarily result in precision.
Yes, it would mean over-resolution over the poles when using a lat-lon grid (other grids like the reduced Gaussian deal with this over-resolution). This has been corrected in the revised version.

Line 327: You have "Yet … yet", maybe start with "Even "?
The sentence has been corrected by removing the first "yet".

Line 340: If any of the critical checks fail (instead of fails)
This has been corrected in the revised version.

Lines 405 and following: this is an important piece of information, and deserves more prominence, as well as discussion. These are very impressive, but are these SYPDs sufficient to enable the DestinE objectives? It would be important to tie these statistics to what is said around line 412. Also, is there scope for even higher speed, or should resources now be dedicated to ensembles?
This throughput is sufficient to achieve the original Climate DT objectives in terms of climate simulations planned so far. Of course, higher throughput is always desirable. Work is ongoing as part of the Climate DT development. Uptake of innovations from a number of research initiatives is taking place, all with the aim to increase the climate model throughput. There are also efforts to increase the capability to make a more intensive use of GPU-based computing platforms (especially with the IFS-based models).
The search for higher speed and the production of ensemble simulations are not conflicting. In fact, ensembles of 10-km simulations have now been performed. We aim to both continue to improve the performance of the models, and carry out small ensembles, making the best use possible of the computing resources granted by EuroHPC.

Lines 442 and following: this is an important part of the methodology, and it would be good to explain in more detail what the technique offers and what its limitations are. For example, moving towards the generation of ensembles, what are the implications of using this technique?
We agree that it is important to offer a bit more detail on the methodology, including its limitations. We have slightly expanded the description of these simulations and the lessons we have recently learned about their pros and cons. Additional references of manuscripts either submitted or recently accepted cannot be included because a doi is not yet available. Ensemble simulations have been used in these publications and the pros and cons of this methodology are discussed with a lot of detail.

Line 542: define data bridge as soon as it is used, possibly around line 355, so that this sentence can be understood.
We have defined the Climate DT boundaries and the relevance of other important DestinE components such as the Data Lake, of which the data bridges are part, and the Service Platform in the revised section 2 to better guide the reader.

Line 794: Synergies require identification of synergies. This seems circular. Can you rephrase?
This part has been rewritten in the revised version of the manuscript.

Reviewer #3

This is one of the most important papers of 2025. It describes the implementation of a critically needed paradigm shift in the way climate model information is produced, disseminated and integrated with decision making. The paper describes the implementation of the workflow required to achieve this, from the execution of climate model experiments, to the handling and serving of the vast amount of data produced to the application of the data in applications. As the DestinE project is in the prototyping phase, there are only a few examples of the latter available at this stage.
The paper is overall well written. It feels long, but each of the innovations described in it deserve to be there. Just like the workflow would collapse by removing one component, so would the paper. While one could always debate the use of particular words here and there, I feel the paper is too important to do so and recommend publishing it as is.
We thank the reviewer for the kind words about the relevance of this manuscript. We have tried to shorten the paper and improve its readability.

I have a few minor comments, which I list below:
1) One of the points made in the paper a few times is that the use of high-resolution global models creates opportunities for a more equitable use of climate information across the world. At the same time, all the examples for applications shown are European. It would be nice to add one or two sentences in the "Looking Forward ..." section on how data consumers in other nations might be able to interact with the system in the future.
We have initially focused on examples of European interest. However, as the reviewer mentions, the Climate DT scope is global and applications beyond Europe can benefit from it. An example that has been worked on in parallel to this manuscript is an analysis of the climate simulations over areas with small islands including the Canaries, La Réunion, and the Caribbean, in a comparison with existing climate projection sources for those areas. As this work is not yet submitted, we have modified sections 6 and 8, which has been largely rewritten following the comments from another reviewer, to make clearer the global scope of some of the applications.
As the reviewer mentions at the beginning, the manuscript does not allow to illustrate all the Climate DT potential applications. One of the objectives of this manuscript is to inspire the community to develop additional examples that can benefit from the equitability principle that guides the Climate DT.

2) While the authors have been very careful in explaining most of the many acronyms used in the manuscript, a couple slipped through the cracks. I found YAC and CSC-IT, but I suspect there are more. Please check the text carefully as the use of acronyms excludes certain groups of readers. especially those who are not part of the project.
Given the large number of acronyms, we have created a table that collects all those used in the manuscript.

3) On line 96 the authors state that the information provided must be "always matching the spatial scale of the decision". That sounds too absolute to be true. Many atmospheric phenomena will be affected by topographic features at scales much smaller than even the relatively high resolution used here. Maybe, "always" should be more like "where scientifically defensible and technically possible".

With this comment we realise that we have not distinguished between climate information and its sources. The sentence refers to climate information, which is a synthesis of material from different sources such as expert views, data (observations, simulations), and literature, among others. The climate information distilled through this process needs to match the spatial scale of the climate-related decision. To include situations in which the decision is not fully identified, as often occurs in climate services, we will substitute “always matching” with “is expected to match”.

4) Line 111: "... with many ..." This does not seem to fit here.
This clause has been removed in the revised manuscript.

5) Line 154: I do not understand how one would configure data consumers. Please clarify.
The sentence has been reformulated to “A digital twin of the climate system targeting adaptation is expected to make use of observations, integrate several climate models to consider uncertainty sources, include applications for climate-sensitive sectors directly connected to the climate models, and provide adequate interfaces to configure the simulations, their output, and the timely interaction with data consumers”. More details about the system configuration have been added in section 3.