EuroCC WP 8.2/8.3/8.5 DK Survey

David Gray Marchant and Tobias Lindstrøm Jensen
david.marchant@di.ku.dk, tlj@its.aau.dk

University of Copenhagen - Aalborg University - June 6, 2022

Contents

1 Executive Summary

As a part of the EuroCC project DTU, AAU and UCPH conducted a survey of HPC resource usage at Danish Universities. The intention was to get insights on the state of HPC, AI and HPDA provision, training offerings and needs within the Danish Universities. This report represents the results from the survey and is a contribution to the overall aim to support that a sufficient number of people are trained for expanding HPC demands in academia and industry.

The survey was conducted in the autumn 2021, and sent to all Danish Universities. 201 questionnaires was answered and returned, even though some was only partly answered. There is a clear overrepresentation of respondents from DTU and science in general, just as other universities only have few or no answers at all. Still the report gives useful knowledge of how the respondents find the supply, content, format and quality of HPC, AI and HPDA provision and training at the Danish universities.

1.1 Key Findings

The main results from the report are:

- The respondents expect increased demand for HPC resources, especially for GPUs.
- Even though many respondents find no significant challenges using HPC, AI and HPDA, many more cannot locate sufficent support or guidance in order to either access or utilise existing resources.
- Most respondents have the right knowledge and skills to participate in offered HPC, AI and HPDA courses
- The respondents prefer shorter and self-paced learning over longer and structured courses
- The available courses seems to favour theoretical aspects over practical, where course participant prefer an equal balance.
- The reports identify that the respondents at universities find shortcomings in access to computer resources and the need for more visibility for both access to HPC resources and courses with a changed format.
- It must be taken into account, that the respondents are spread across all sciences and largely already users of HPC, AI and HPDA.
- It is expected that there will be an increased demand for HPC, AI and HPDA within other disciplines that may not be reflected in the respondents.

2 Introduction

The European High Performance Computing Joint Undertaking (in short EuroHPC JU) aims—among other things— at building an European network of 33 National HPC Competence Centres (NCC) via the EuroCC project. In Denmark, DeiC is coordinating this work, in seven workgroups. The Danish group WT 8.2 "Training and skills development" is to facilitate effective training and skills development in the use of HPC, AI and HPDA in Denmark.

Throughout this report we will refer to these three areas (HPC, AI and HPDA) collectively as PDC. This is done for ease of reference as in many contexts their definitions each overlap, or are used interchangably by users. Specific terms, such as HPC, will still be used in existing definitons and where they are directly important to interpreting the results of individual questions.

To get a better understanding of the current state, in the summer 2021 the group decided to form a survey jointly with the EuroCC WT 8.3 "Technology Transfer & Business Development" and EuroCC WT 8.5 "Mapping of HPC/HPDA/AI Technical Competences" groups.

In both the short and long term, we expect to see an increased interest in PDC throughout scientifc research. This is both motivated and supported by initiatives such as the EuroJU/EuroHPC/EuroCC projects, as well as the development of core undergraduate topics in PDC. This is not limited to Computer Science, but is also present in a variety of Physics, Chemistry, Biology, Engineering, Business and other research areas. An initial statement for PDC is shown in Fig. 2, and can be viewed as an extension of "The

The principal goal of this task is to oversee and coordinate Denmark's national and local training and skills development within HPC. This organization shall keep track of and actively contribute to develop:

- *Educational programs, including graduate and PhD-courses, courses for beginners and advanced HPC users etc.*
- *Industrial outreach activities*

Figure 1: WP8.2 description

Carpentries" with an HPC branch [1], and local initiatives (Scandinavian Coderefenery [2], UK Software Sustainability Institute[3]).

"PDC is not just an integral part of the work of computing professionals who explicitly design systems that exploit concurrency to achieve performance; it is also relevant to developers and users of applications that hide much of the complexity of harnessing PDC technology. These implicit consumers of PDC may include developers with applications that interface with everyday tools or libraries such as collaborative environments, productivity tools, and multimedia applications that utilize local and/or remote PDC technology implemented below their visibility threshold. The penetration of PDC into the daily lives of both "explicit" and "implicit" users has made it imperative that all computing professionals be able to understand its scope, effectiveness, efficiency, and reliability."

Figure 2: Statement by an NSF/IEEE-TCPP working group for topics in parallel and distributed computing (PDC) [1].

Based on the above task, and the idea of a broader penetration of PDC, the group worked with the following five questions:

- Do the users have sufficient skills and competences for their daily work relating to PDC?
- Are the barriers for entry for those looking to start using PDC too high?
- Can users locate sufficient support and guidance in order to successfully utilise PDC?
- What is the perceived communication, accessibility and quality of the current PDC competence development, and do these develop further possibilities?
- Are there significant gaps in the current offerings that WT 8.2 can address?

3 Methodology

Based on the previous five questions, a survey was constructed by EuroCC WT 8.2 during the summer of 2021. The questionnaire was send out via DeiC[4], an organisation for fostering networking between Danish universities and PDC resources. The survey was sent to the front-office or equivalant of each Danish University on August 16th with a request to forward the survey to any and all relevent people within their respective organisation. A short letter justifying and explaining the survey was attached from Allan Have Sørensen, and is shown in Appendix G. The survey itself was hosted at SurveyXact, an online platform for conducting and managing surveys. Respondents were given the deadline of September 15th 2021 to complete their submissions.

[1] `https://www.hpc-carpentry.org/`
[2] `coderefinery.org`
[3] `software.ac.uk`
[4] `https://www.deic.dk/da/node/1978`

4 Initial comments on survey

The survey was forward via DeiC to the front-office of each university. After the initial survey closing we observed a low response rate from SDU and AU, and decided to contact key persons here to see if it is possible to collect information from these as well. The suvey is open to anyone with the link, so could also included false responses and multi-responses from a single individual.

A few initial observations on the survey.

- We received 143 complete and 58 partially complete responses. Due to the method by which the survey was handed out, we do not know how many individual respondents were contacted so we cannot calculate a response rate.
- There is an overrepresentation of respondents working at DTU. We should take this into account.
- Phsyical Science may be overrepresented in the results as the front offices may have only forwarded the surveys to those they would expect to use PDC, such as Computer Science or Physics departments rather than Philosophy and Literature.
- Many had used the free text fields to comment on the questions and the current state. We should take these into account in the analysis.

5 Respondents background

We will now consider the makeup of the respondents by:

- Institute/company
- Position
- Field of interest

5.1 Institute/company

When looking closer at the distribution of responses in Fig. 3 and 4, it is clear that we observe a large amount of responses from DTU. Two universities, AU and SDU have only a handful of respondents, while RUC and ITU have no reporting respondents at all. This could be a sign that most HPC is not needed at these universities, but this is unlikely. It could also be a sign that outside of DTU few people recognised EuroCC or those attached to this project and so dismissed the survey out of hand. As sevearl people in the project work in or closely with DTU, this could explain the much higher response from there.

This will limit the generality of the conclusions that are possible to draw from this survey, though when breaking down the makeup of respondents we can see that DTU respondents are broadly representative of the general makeup of other institutions, with a heavier focus on materials and other physical sciences as would be expected. We also observe that some groups of respondents are quite low, e.g., computer science from AU. It is also of note that several respondents replied with 'Public Instituition', by which it is assumed that they mean a university. As this is not known for sure they have been kept seperate from from those who explicitly answered 'University'. Other institutions listed are NBI, which is part of KU, as well as Computerome which is part of DTU and KU. Finally is Demant, a large Danish medical company. As this survey was only handed out directly to universities it is unclear exactly how they accessed it, but presumably they have some ties to one of them.

5.2 Job/title/position

We now consider the question "Your current job title (select the most suitable)" with the distribution shown in Fig. 5. We mainly observe senior and junior university staff categories, with only a few respondents in the remaining categories. This shows that these results and any conclusions from them are heavily weighted towards academia, and its direct interests, and does not include any significant input from industry. It is also of note that the single largest category by some margin are the junior researchers. This is somewhat unsurprising as it is they who would typically be doing most of the actual implementation in order to use HPC resources.

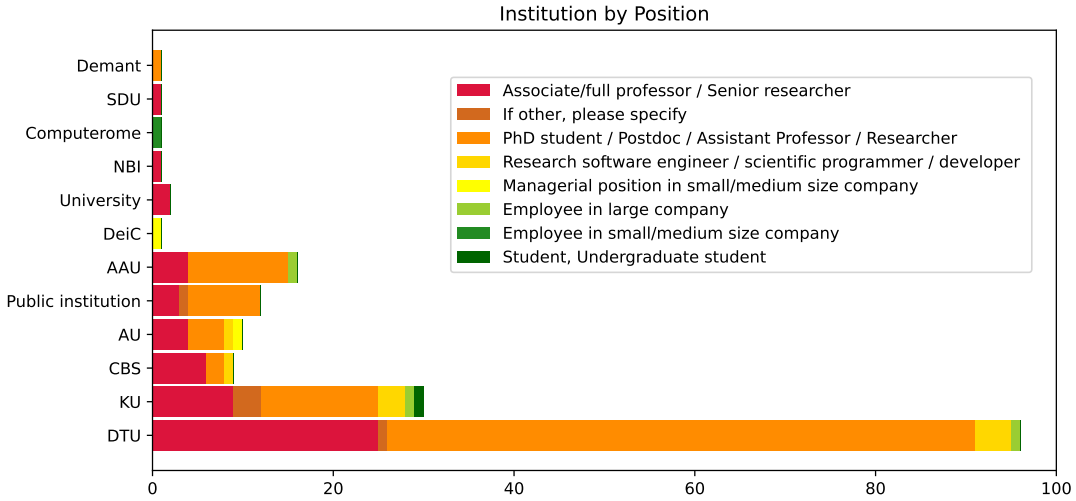

Figure 3: Distribution for question 1, number of respondents broken down according to the position of the respondent.

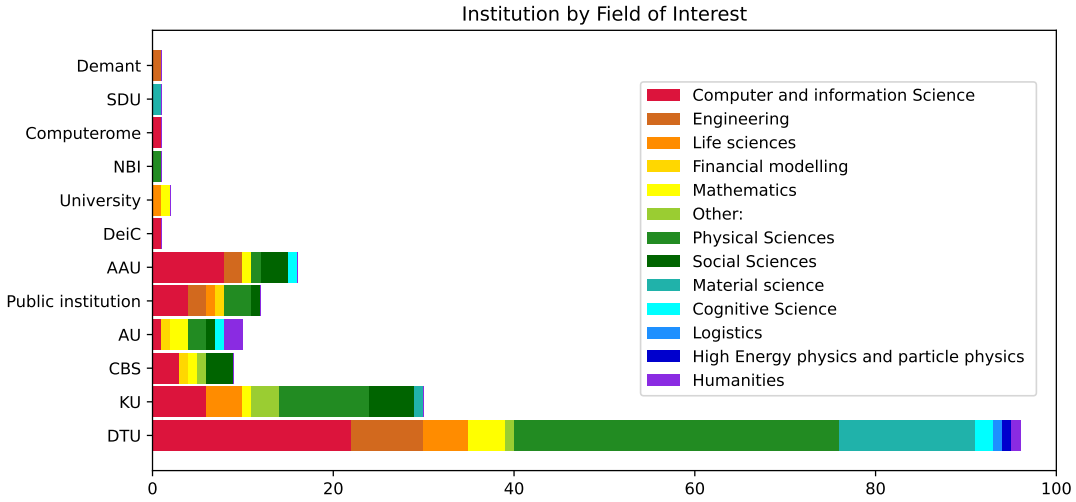

Figure 4: Distribution for question 1, number of respondents broken down according to the field of interest of the respondent.

5.3 Field of interest

We will now look closer to the responses broken down by field of interest shown in Fig. 6. We observe a large number of responses from "Computer and information Science" and "Physical Science" with a mid group of "Engineering", "Life Science", "Social Science" and "Material Science". The remaining groups are getting too small to conclude anything of statistical significance for these fields. Note that DTU dominates the responses for Materials Science and Phsyiscal Sciences, as would be expected as a large, technically focused university. We observe a low number of Other/NaN indicating the available choices was sufficient in most cases.

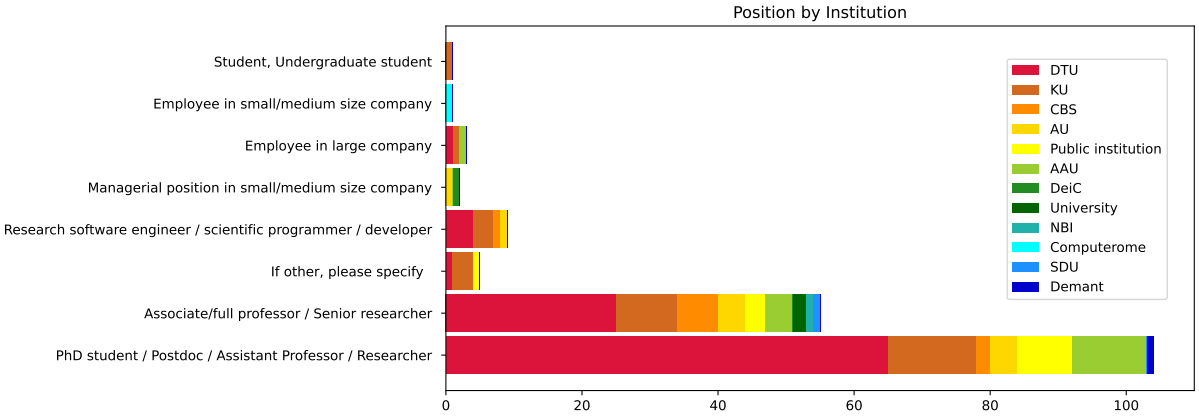

Figure 5: Distribution for question 3: Amount of respondents with job title, broken down according to hosting institution.

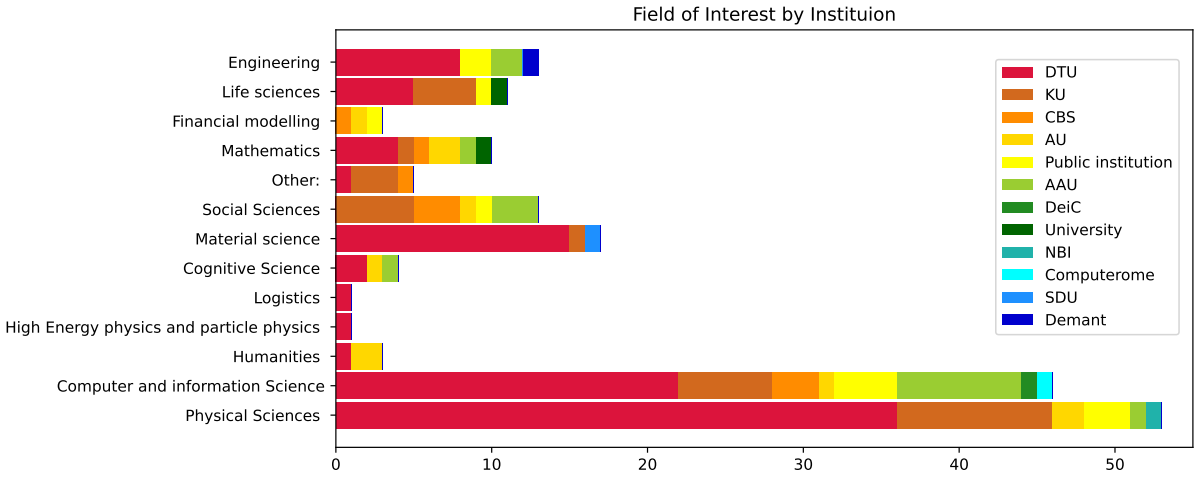

Figure 6: Distribution for question 4: Amount of respondents in each field of interest, broken down according to hosting institution.

6 Usage and challenges

6.1 Computational considerations

Within questions 7 and 8, respondents are asked about their future expectations regarding HPC. The results of these questions can be seen in Fig. 7. From these we find that the overwhelming majority of respondents expect to use increased computing resources, but that many of them do not have any real plan on how to do so.

It is theorised by WT 8.2 that the growth in HPC usage in future will primarily come from outside of Computer Science, as other fields of study require more HPC resources in order to perform their analysis or experiments. This is consistent with our findings here as we have quite a cross section of fields within the respondents, and yet almost all report benefiting from more HPC resources.

6.2 Experience

Question 9 asked respondents how much experience they had with either HPC, AI/ML, or HPDA. The results are presented in Fig. 8, though perhaps the most interesting thing to take from these is that all fields have much more experience with HPC than that with either AI/ML or HPDA. This could perhaps be an indication that these are less generally applicable to problems, or that they are newer

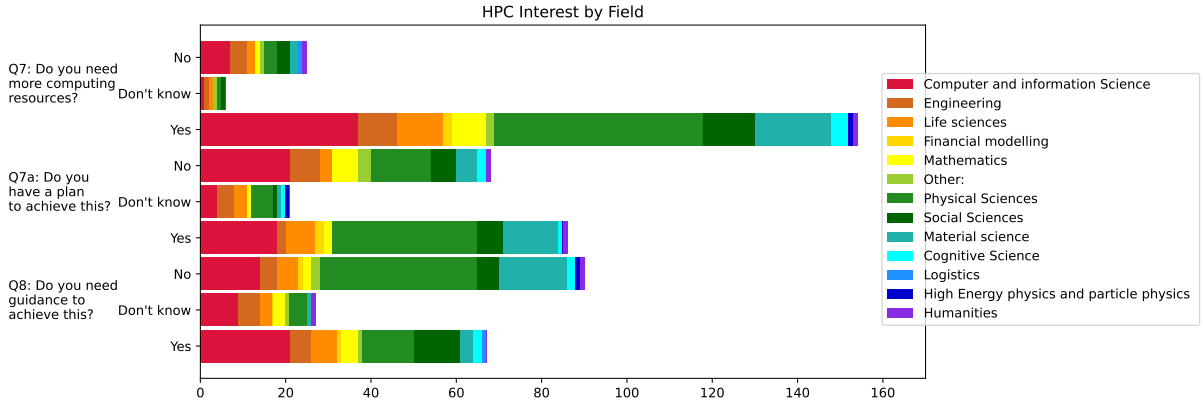

Figure 7: Distribution for question 7 and 8: Amount of respondents with each computational considerations. Note that for brevity the questions have been abridged. Full text is available from Appendix H

fields and are not yet being utilised to their full potential. However, it could also be simply that HPC is a more recognised term and so non-expert respondents report HPC where AI/ML or HPDA could also be applicable.

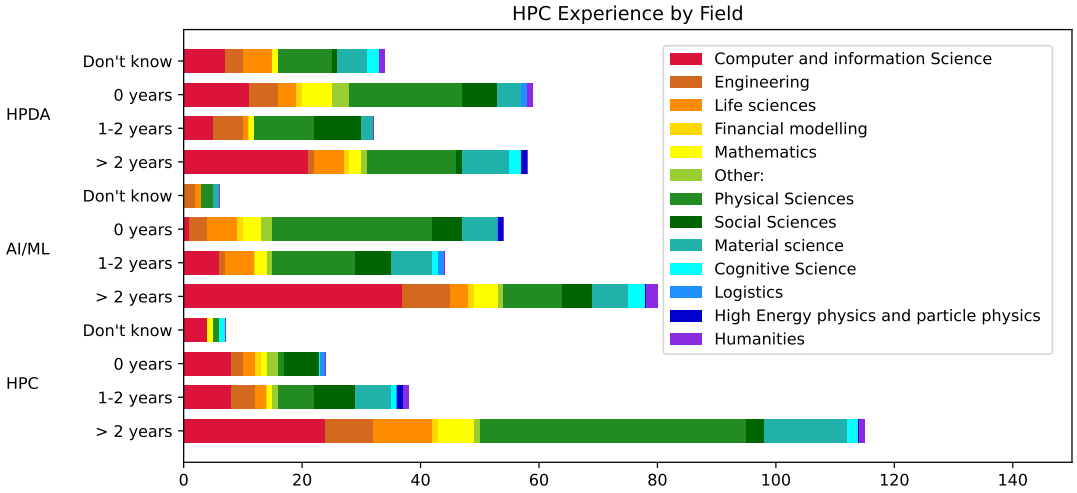

Figure 8: Distribution for question 9: Amounts of respondents with each experience level.

6.3 Daily interaction

Question 10 groups respondents by their daily interaction: "What best describes your daily usage and interaction with computer software?" shown in Fig. 9. We observe a broad spread on the mentioned categories with a higher response for low-level interaction. By field, we also observe broad response for all fields in almost all interaction categories. From this we can conclude that many respondents are implementing their own solutions to problems. This could indicate that either individual problems require individual solutions, or possibly that more general solutions/frameworks need to be made available to researchers. This may be worth future investigation.

6.4 Features

Question 12 tries to address new features with the questions "Which features would you like to add to your application? Select all that apply (leave empty if you don't know)". The distribution is shown in

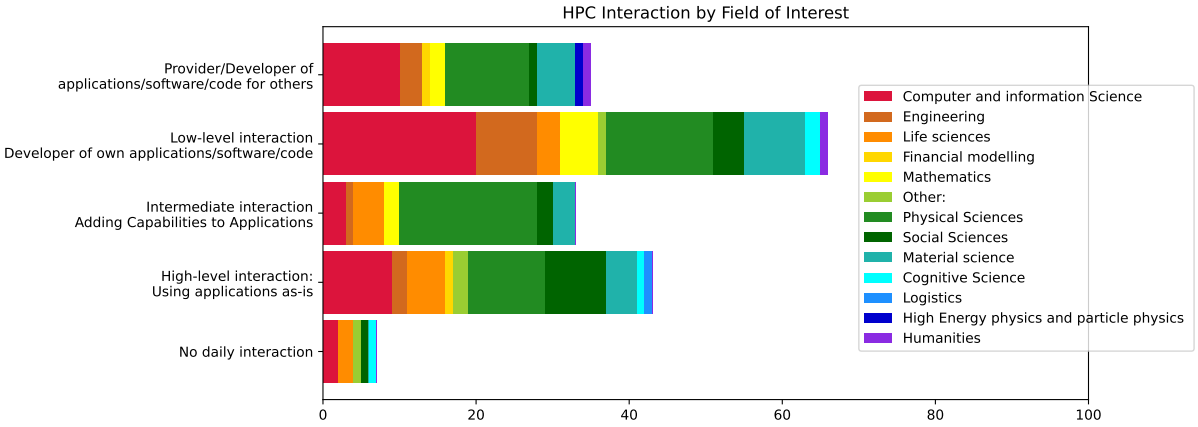

Figure 9: Distribution for question 10: Amount of respondents per daily PDC interaction, broken down by field of interest.

Fig. 10. It is interesting to note, that there is a big mid-field with between 20-40 responses. This seems to indicate a wide field of users interesting in different features. When it comes to accelerators, its clear that GPU computing is much more sought after compared to FGPA. It was assumed that different subject areas would have differing requirements, but once broken down by each respondents field of interest it is apparent that this is not really the case. From this we can conclude that any resource provision should accommodate all fields, rather than say focusing all GPUs towards Material Sciences.

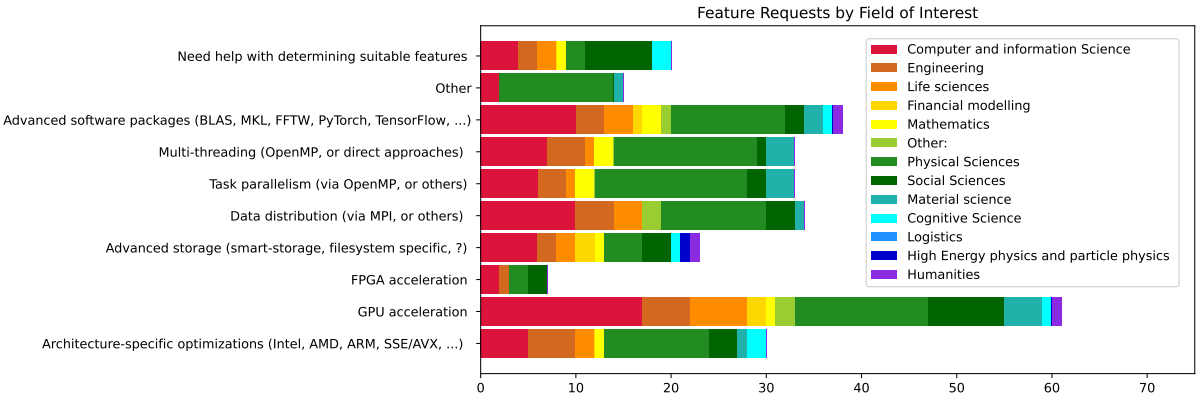

Figure 10: Distribution for question 12: Amount of respondents requesting features to add, broken down by Field of Interest.

The free text replies are below, all of which are from individuals at DTU. Note the dominating request for more compute resources, in-particular CPU:

DTU More CPU compute nodes are always welcome!
DTU I am fine with the features I have
DTU Increase raw computational resources
DTU Increase CPU capabilities
DTU I'll need more GPUs available, because there are periods where none is available
DTU High performance visualization
DTU CPU
DTU more CPUs
DTU cpu needed
DTU I have what I need already
DTU Maybe more public-accessible CPUs please
DTU More CPUs
DTU CPU
DTU more CPUs are in demand
DTU I would personally like to have more CPUs

These results are consistent with the responses to Q13 as discussed in the following section, where the main identified hurdle from respondents is the difficulty in acquiring more compute resources.

6.5 Challenges

Question 13 is on individual and organisational challenges with the question: "What are your primary needs and challenges for yourself or organization? Select all that apply.". The response is shown in Fig. 25.

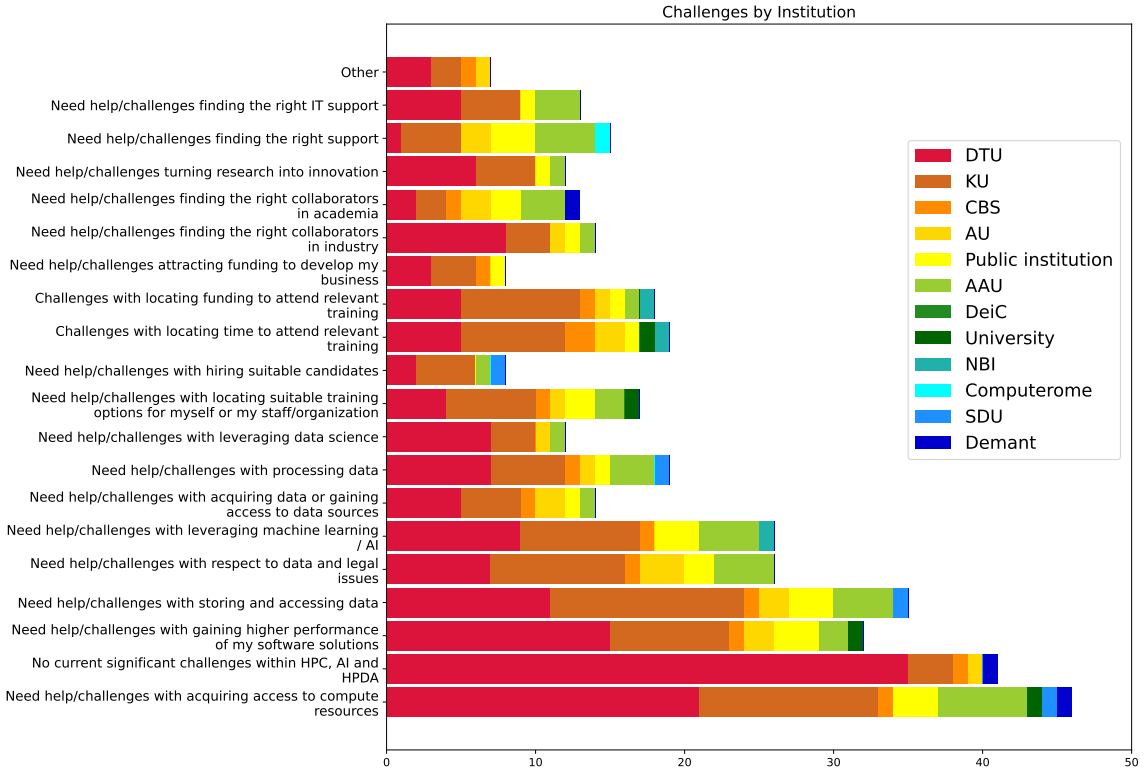

Figure 11: Response for question 13: Amount of respondents with needs and challenges, broken down according to Institution. A larger version is available in Appendix

Considering Fig. 25, the challenge with most responses seems is "Need help/challenges with acquiring access to compute resources" (35%), with top responses also on storing and accessing data (26%), performance issues (23%), leveraging machine learning (19%) and data/legal issues (18%). Direct issues in connection to competences are interesting at a lower level, including locating suitable training (13%), time to attend (12%), funding (12%) and hiring suitable candidates (7%). Interestingly the second highest response is "I find that I have sufficient competences and skills to conduct my current work within HPC, AI and HPDA (no significant challenges for the time being)" (29%).

We note that all selections have been used with the lowest at 8 responses and only a few other:

DTU Need to allocate more resources for HPC compute nodes (CPU).
DTU Lack of computing time
CBS ?
DTU Setting up more GPUs
KU Simple introduction to students to use Computerome or similar resources
AU I am on a very basic level, so more training would be desired.
KU I need more and clearer information about available resources

The free text replies are shown in Appendix D, where a wide range of issues are raised like support, technical issues, man power, funding and lack of computing resources.

The top response "Need help/challenges with acquiring access to compute resources" indicates that it is important to supply resources, especially in non-DTU and non-physical science settings. Many local and field specific challenges could lay behind such issues. For instance, is the barrier an actual lack of resources in certain fields and institutions, or is it a lack of information of existing computing resources?

One alternative interpretation of these results could also be that finding resources is the first step in using them, and so will be over-represented in the results. This may be the first hurdle, and so any respondent who does not know anything will fail here before they even get to the next step and realise

they don't know that either. For example, a respondent who cannot find a resource to use, may not yet realise they also don't know to use a resource were they able to find it.

"Need help/challenges with gaining higher performance of my software solutions" is an interesting and high ranking challenge. The first question is whose task is it? If it runs fast-enough for a researcher to obtain their results, then only the administrators have an interest in terms of saving compute resources for others. On other hand, skills in performance optimization are often seen as a specialization requiring in-deep hardware/software understanding. Perhaps it makes more sense that only some people learn this skill as a specialization, and not as a widely used competence.

The second highest response is on "no significant challenges for the time being". A dive into the data indicate that of the 43 responses in this category, 34 comes from DTU, with 25 from Physical Sciences and 13 from Material Sciences. That is, the response are dominated by very specific institutions and/or fields of interest and should not in general be seen as widely spread response. It is also likely that these responses comes from areas with well-defined problems, maybe even such that hardware and software setup can be co-designed for particular problems. This makes the problems easier to work with and more similar when going from case to case.

6.6 Support

To get more insight into the type of requested PDC support, as well as that which is currently provided, we can combine questions 14 and 15. This is so we can investigate what respondents want, and if there is a clear mismatch in the current provision. Fig. 12 shows the collated results of the questions "Which of these support actions would you be interested in?" (question 14) and "In which of these areas, if any, do you provide support to your own and/or another organization?" (question 15). We do not believe that exact numbers can be directly compared with each other, as a single provider can service many interested parties. We do see a wide range of interest and provision, so in principle at least there exists the possibility that all interests are provided for.

- Question 14 other free text:

SDU	Support from SDU IT services. To solve problems related to SDU IT policies develoepd for word users, and not always useful for HPC users.
CBS	?
DTU	Providing more computational resources

- Question 15 other free text:

SDU	I offer support to my group (Ph.D. + project students) in using my computational facilities, storage needs, and post-processing of data.
Public institution	I don't provide support as such. But I collaborated supervise on many data intensive projects.
CBS	none
NaN	No support given
DTU	nothing

7 Courses

7.1 Traning interest

Fig. 13 shows answers for question 16, broken down according to the position of the respondent. Question 16 reads "For which of these subjects, if any, would you be interested in training for you or your staff? Select all that apply.". Of the 201 surveys recieved, 75 did not answer this question. There is an option for respondents to say that they have no interest in additional training, though this is only one of twenty eight answers so may have been easily missed. Similar to other questions, we observe a broad spectrum of responses indicating a wide usage and interest across the suggested areas.

The free text field for other contained:

DTU	I?m not interested in additional training. I would prefer more HPC hardware.
DTU	I?m not interested in additional training.
CBS	?
DTU	I would just like access to more computational resources
Public institution	Good and well thought approaches for storage and curation of large amounts of data.
SDU	Essentially I have no time for training, but is interested in developing technologies such as DPC++, cuda etc.
DTU	Not interested in external training.
DTU	No additional training
DTU	Courses for PhD students about visualization of data (ParaView)
DTU	I'm not interested in additional training.
DTU	Training is well provided from my university

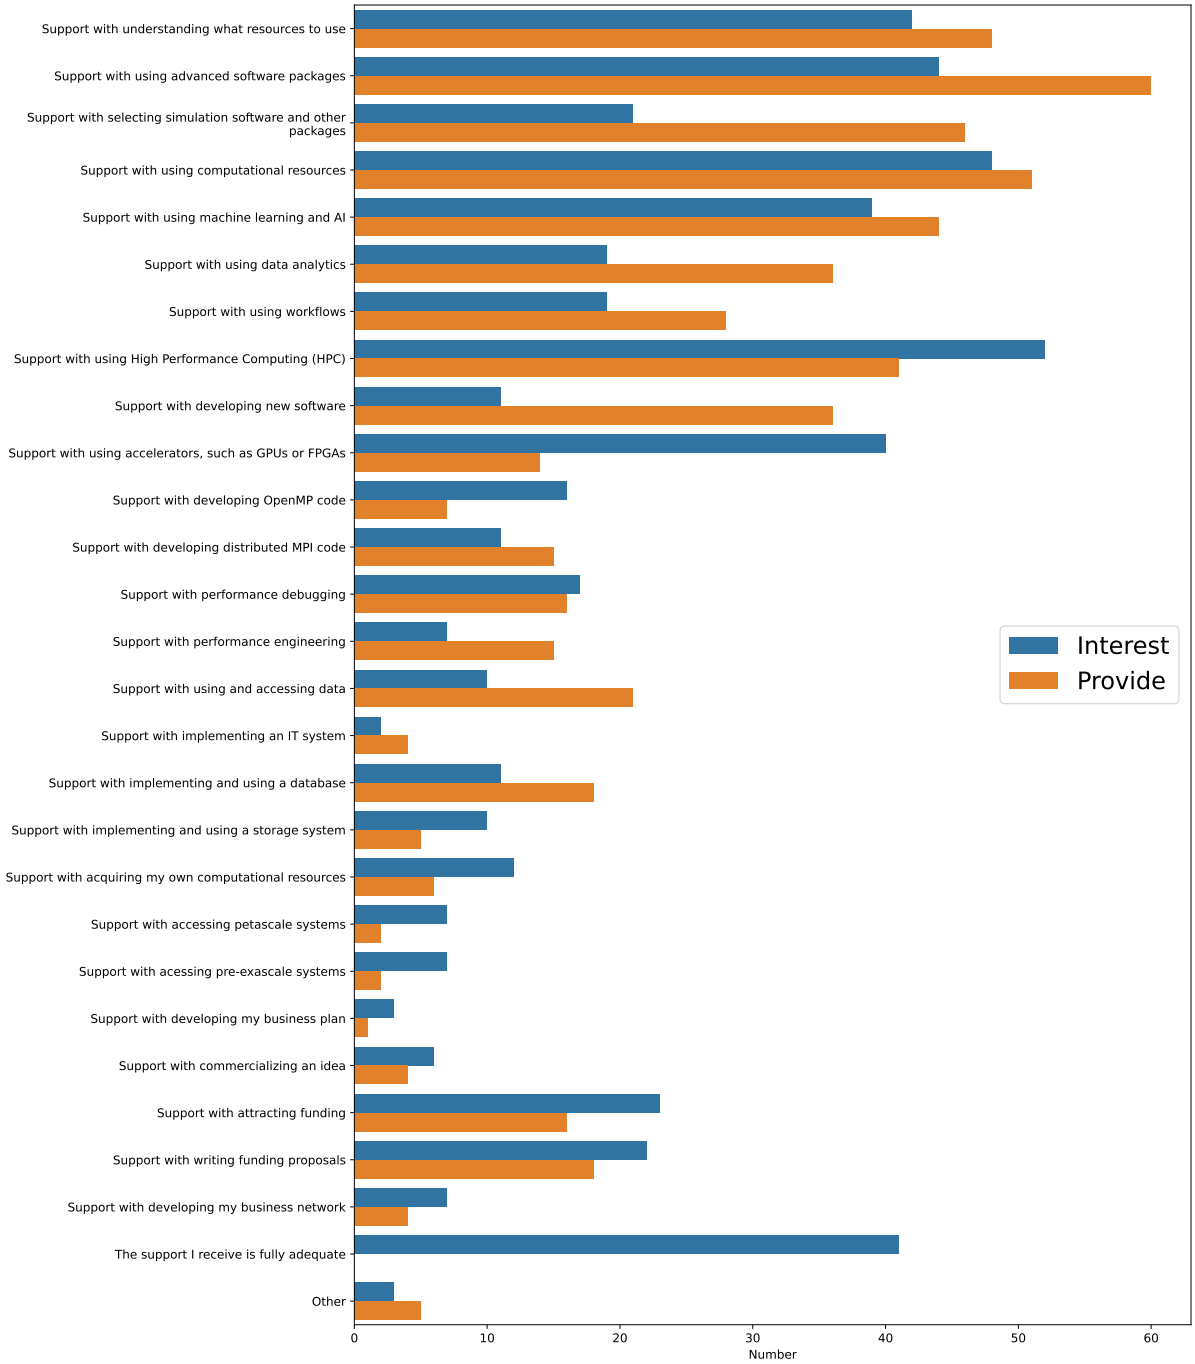

Figure 12: Response for question 14 and 15: Amount of respondents either interested in or providing support

DTU	I do not need additional training
DTU	I do not need additional training
DTU	No need for additional training
NaN	Don't know now – it's project dependent
AU	developing software for users without programming experience
DTU	Do not require additional training
DTU	I'll be fine without training via formalized outside contacts.

The main content of these indicates no interest in additional training by the respondents.

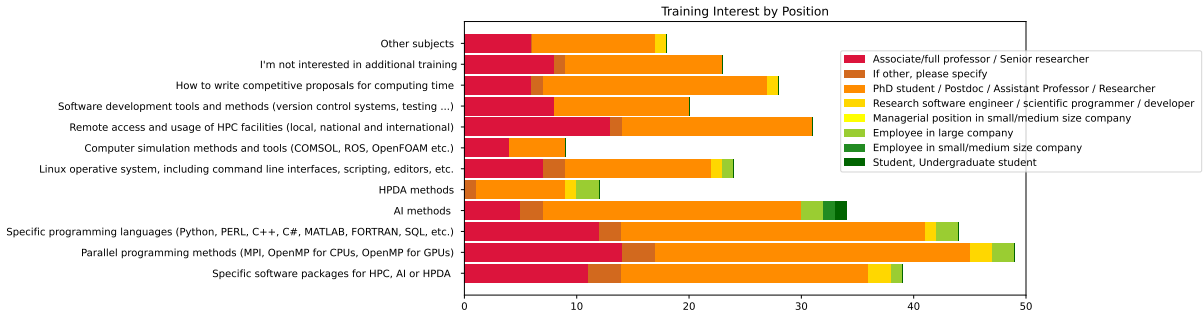

Figure 13: Response for question 16: Amount of respondents with each training interest, broken down by position

7.2 Preferred type of training

The matter of preferred training type for competence development is addressed by "Which of these training events, if any, do you find most useful? Select all that apply." shown in Fig. 14. We observe highest response for workshop style, followed up by self-learning. In total, 66 did not respond to this question. From a competence development perspective, it is interesting to note such a wide support for self-learning. Based on the 8.2 task description in Fig. 1, it is unclear how we can support and develop self-learning in the given frame.

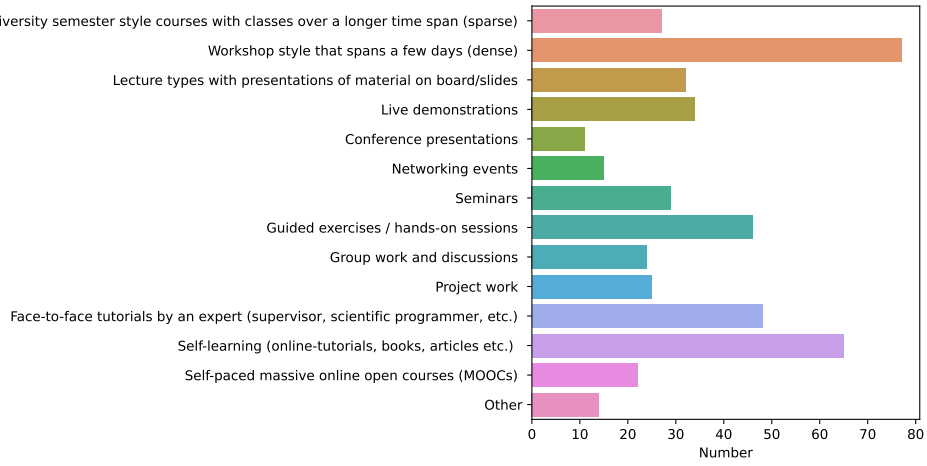

Figure 14: Response for question 17: Amount of respondents preferring each training type

There is a smaller interest in classical university style courses with classes spread over a longer time span (sparse). Face-to-face expert interaction enjoys a sizable level of support, perhaps due to it being effective in locating problems fast. However, this is in general quite an expensive solution, as its one problem at a time. This makes it difficult to scale. On the contrary, MOOCs scale very well, but here we observe a somehow lower interest in this type of training.

We continue with the question "19. Do you prefer in-person or online training events?" shown in Fig. 15, with 56 non-answers. Interestingly we observe a lot of support for mixed approaches, followed by interest of in-person over online events. This is contrary to the observations reported in the Swedish EuroCC National Competence Center report [3], where they observed a larger support for online events.

7.3 Event length

Question 18 "How many work-hours or work-days do you think is optimal for non-self-paced training events?" is shown in Fig. 16 and 17. Although both responses have a few higher time frame responses,

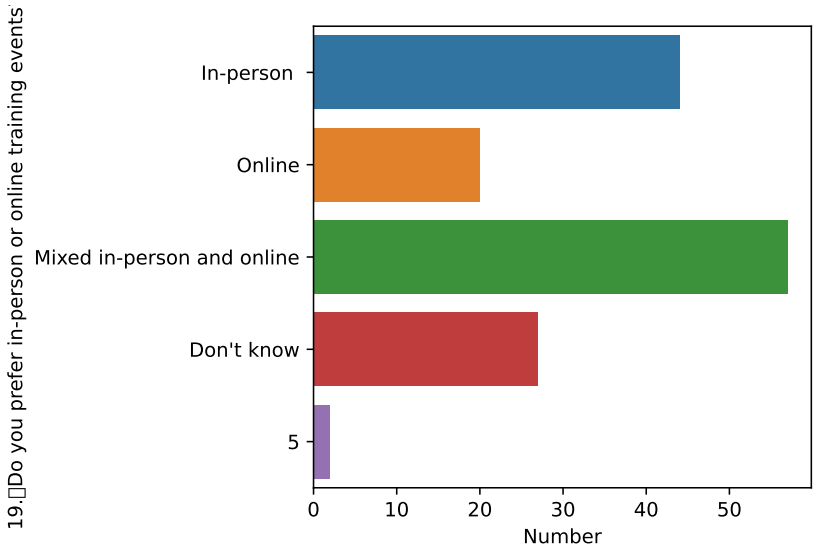

Figure 15: Response for question 19: Amount of respondents preferring either online or inperson training

both show a general preference for shorter training courses ranging from a few hours to a few days. This perhaps reflects that researchers do not wish to dedicate months of intense study to the field of HPC in order to get their research, but are simply after relatively quick and easy ways of getting their processing onto large compute resources.

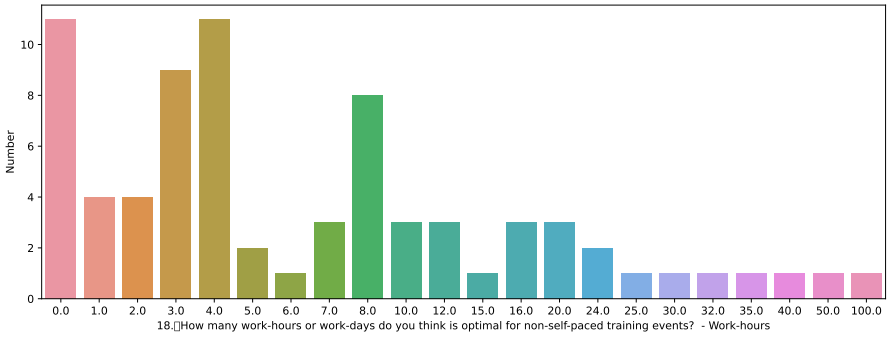

Figure 16: Response for question 18: Amount of respondents preferring each event length in hours.

7.4 Evaluation of last course

Question 21 is divided into four sub-questions. The first three are a five step grading (poor,—,average,—-, good) on "If you previously have taken courses on programming languages, or courses regarding HPC, AI, and HPDA,...". The three questions relates to overall experience, form and presentation, and if the course lived up to its description and the response is show in Fig. 18.

As is shown in Fig. 18, only 93-97 respondents answered these first three questions. Another key observation is that even when taking the least satisfied answer (lived up to its description), more than

$$\frac{94-8}{94} \approx 91\% \tag{1}$$

of respondents have rated their last course at average or above across all three questions. Its also interesting to note that more than half of responses are in the average and good category compared to

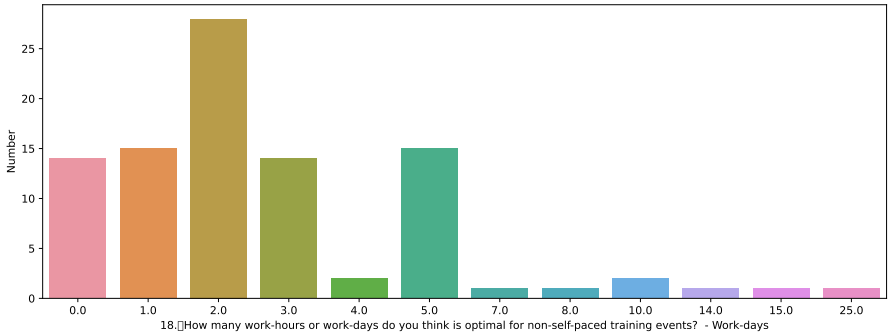

Figure 17: Response for question 18: Amount of respondents preferring each event length in days.

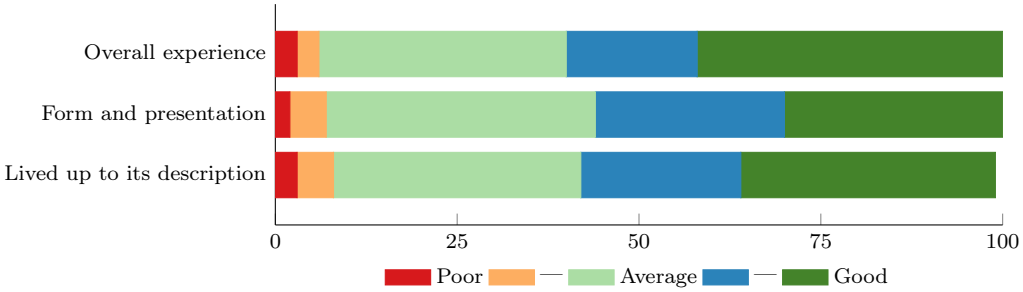

Figure 18: Amount of respondents rating each of their last PDC course.

the category in-between. It is suspected that this just illustrates that people will not pick an unlabelled category, and we should have labelled them something like 'very poor, poor, average, good, very good', or should have stuck to just three options.

Question 21: "Consider the same course you last attended. What issues did you encounter for this course. Select all that apply" responses are show in Fig. 19, with 3 responding in free text

Demant	The particular course was poorly planned and executed, compared to other courses with the same format. This is not my general experience.
DTU	Not good enough access to computational resources
DTU	Was okay, but I didn't really need to know those things. (Did need the ECTS points though)

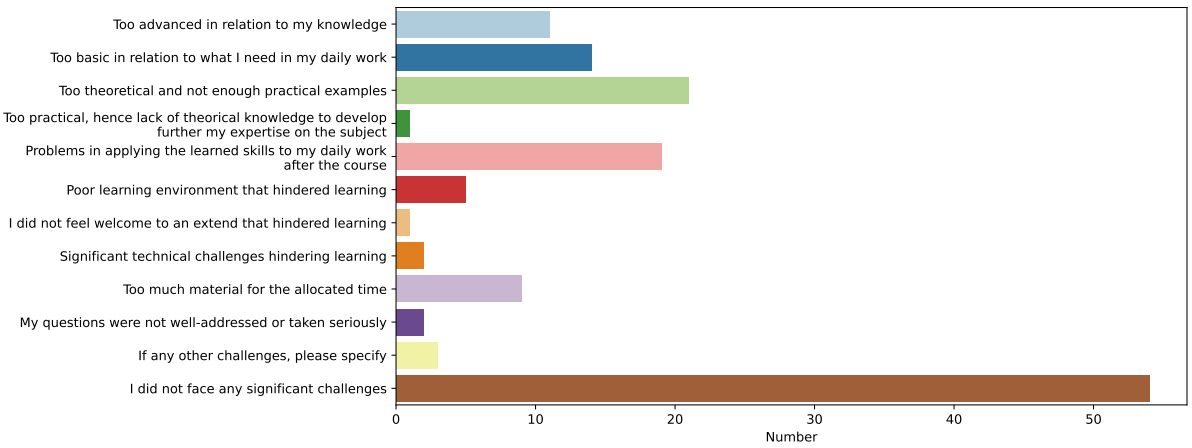

Figure 19: Response for question 21: Amount of respondents encountering issues in their last PDC course

From Fig. 19, the first impression is that respondents did not face any significant challenges at (54),

with follow-ups on "Too theoretical and not enough practical examples" (21) and "Problems in applying the learned skills to my daily work after the course" (19). In contrary, its interesting to note that more respondents replied "Too basic in relation to what I need in my daily work" (14) as opposed to "Too advanced in relation to my knowledge" (11).

Seen from a balancing perspective, note that overwhelming more responded "Too theoretical..." compared to "Too practical...". This may show that educators should, in general, pursue a more practical approach as respondents are not necessarily interested in what they cannot immediately apply to their research. Such an approach might also address the issue of "Problems in applying the learned skills to my daily work after the course". This seems to represent a classical pedagogical challenge, where it is not possible for the organiser to present examples based in each individual attendees case, and that each participant then, at-least to some extend, need to map the learned material to their particular case at hand. Reflecting on this in combination with the previous question about ideal course length, this may be a function of respondents wanting quick solutions to potential complex problems. By this is meant that HPC is a highly specialised subject that other fields will need some awareness of to make use of large compute resources. However, without spending large amounts of time learning the ins and outs of HPC, it may not be obvious how theoretical examples could be applied. Therefore, we might also conclude that conversely, training needs more theoretical parts so that respondents can better apply techniques more generally to their research.

Social course issues such "Poor learning environment that hindered learning", "I did not feel welcome to an extend that hindered learning", and "My questions were not well-addressed or taken seriously" accounted for $5 + 1 + 2 = 8$ responses out of 100 responses to this question. Addressing these types of issues in computing is now more common and visible in e.g. The Carpentries Code-of-Conduct[5].

In Fig. 20 we break the issues based on institution. From Fig. 20 its clear that the distributions are somehow similar, but with a larger group of "I did not face any significant challenges" at DTU and the rest, compared to KU.

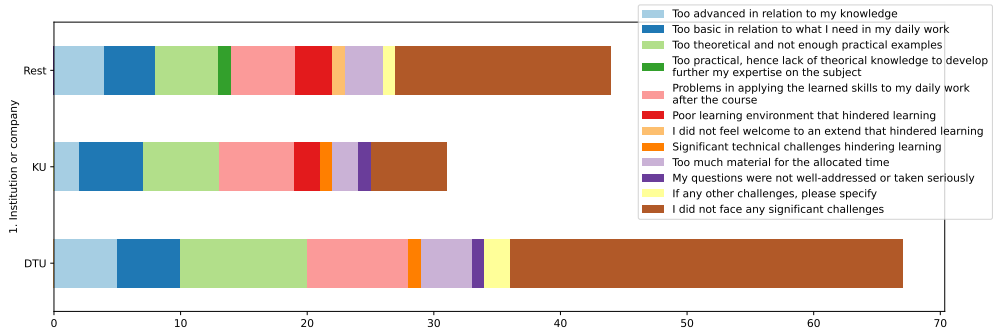

Figure 20: Response for question 21, Amount of respondents with issues in the last attended course, broken down by grouped institutions.

If we instead break the issues down according to job title, we can see some more significant differences as shown in Fig. 21. Here we observe a significant difference in the relative responses between the junior and senior staff with approximately 66% (16/24) of senior staff replying "I did not face any significant challenges", compared to 32% (35/104) of junior staff. A significant number, 18% (19/104), of junior staffs reported "Too theoretical and not enough practical examples", whilst no senior staff reported any such issue. This is perhaps unsurprising considering the expected experience of the two position categories , and the type of work each does, with senior staff on the whole being far more theory and big-picture based. Furthermore, it is likely that junior and senior staff are attending the same courses. So, the additional experience level of senior staff makes them a better match for these courses compared to junior staff, *e.g.*, they have better coping abilities. At the same time, junior staff is likely more in need of adapting the learned skills into their daily work. In particular, looking for ways to apply the learned skills to their particular project.

[5] `https://docs.carpentries.org/topic_folders/policies/code-of-conduct.html`

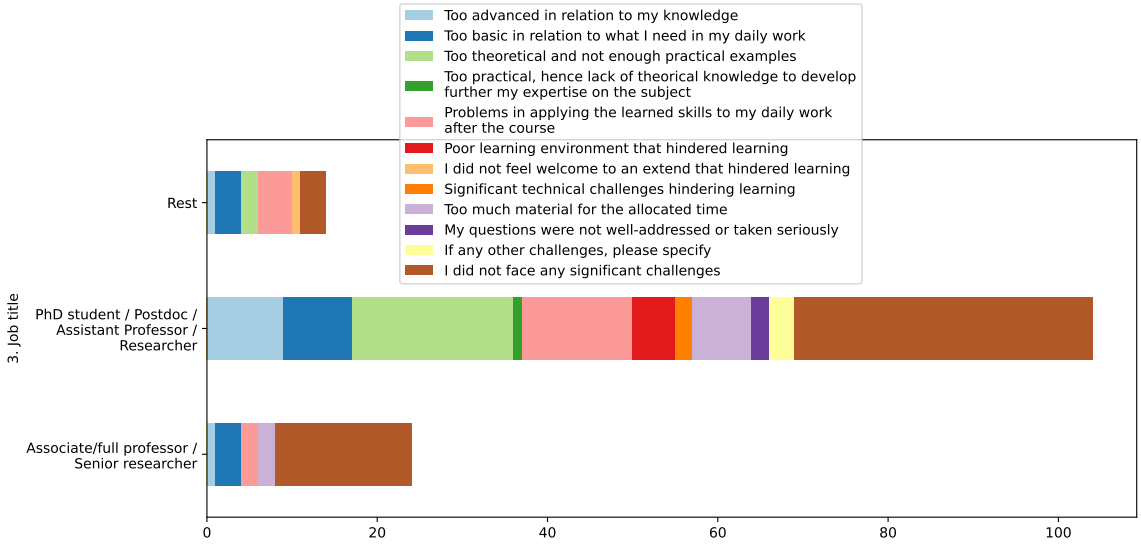

Figure 21: Response for question 21, Amount of respondents with difficulties, broken down by job titles

8 Additional information

Question 22: "In relation to your training needs and/or training experience (for example from your teaching experience), please feel free to write suggestions, ideas, criticism etc. that could contribute to strengthening the knowledge and use of HPC, AI, and HPDA in Denmark" offers some additional insight into the field with a total of 22 usable comments - see section E. The feedback here is very diverse but a summary:

- Diverse feedback with response stating there is a sufficient background in existing university courses to request for additional training - in particular towards master, PhD and postdocs. One points out just to move on at ones own pace.
- Linux is difficult/enforce the use. HPC environments is traditionally Linux environments and this offers some challenges for non-Linux users when first entering this world.
- Training from universities could be workshops, code camps or hackathons. Training on particular "projects" or "software" are mentioned.
- Interesting, its noted that "The number and coverage of HPC courses in Denmark is quite sparse, and no central overview exists".
- Difference between users and instructors, like a computer scientist teaching a group of perhaps non-IT engineers on scientific programming.
- Responds also includes comments on computational resources, and the use of external providers for e.g. handling sensitive data.

Question 23: "In relation to use of HPC, AI, and HPDA usage, please write any additional information about your current needs and what you perceive to be the current state of offers in Denmark" asks for additional free-text information. We received 35 replies (excluding no comment N/A) etc. This is a large fraction of the number of responses indicating a group of people interested in presenting their view of HPC in Denmark. The comments are given in Sec. D. Again the feedback is very diverse, but in summary most feedback is on the themes of "more compute" like CPU, GPU, HPC, funding. If there is one focus it is that "more CPU" is a particularly strong and consistent theme. Another frequent point that was brought up is relating to access and policies. Respondents are often unsure as to what existing resources there are, and what can they access via existing arrangements, such as their university having access to existing HPC centres or not.

9 Conclusions and Recommendations

As discussed in section 2, the creation of this survey was motivated by a desire to answer five questions. Having now reviewed the results of the survey we will consider them in direct relation to these initial questions.

9.1 Do the users have sufficient skills and competences for their daily work relating to HPC?

In general, we observe in the questionnaire that many and a wide range of challenges are raised, but also a significant number of respondents reporting no significant challenges for the time being. This can be found both in respondents ongoing work, as evidenced by Fig. 25 as well as in HPC training as evidenced by Fig. 19. These responses with few problems are mostly coming from DTU, which has a relatively long tradition of hosting dedicated HPC hardware and making it available to researchers.

For non-DTU respondents there is a higher instance of problems, with most reporting their primary challenge being getting access to resources rather than specific skills they lack. This could indicate that respondents for the most part have sufficient skills and competences to run their daily work on HPC resources. However, it should be noted that a not insignificant number have reported issues relating to a lack of skills. Therefore, a tentative conclusion is that some respondents lack skills and competencies, but no particular single gap was observed (other than a lack of ressources). We instead observed a range of issues as evidenced by Fig. 25.

9.2 Are the barriers for entry for those looking to start using PDC too high?

There are a number of respondents reporting lack of time and funding to attend otherwise relevant training, shown in Fig. 25, and respondents reporting no interest in additional training in Fig. 13. But overall these represent a minority response.

It is also worth noting in Figure 25 that the main difficulty reported is that respondents cannot get access to PDC resources, rather than gaining the knowledge or such to access these resources. This may support the idea that there isn't a great barrier of entry for PDC. However, this could simply be respondents stumbling at the first hurdle and reporting the first problem they encounter, even though there would be more siginifcant barriers if they were able to access some resource. More research in this area would be needed before any real conclusion could be formed for this point.

9.3 Can users locate sufficient support and guidance in order to successfully utilise HPC?

It is very notable, that getting access to compute resources is a major challenge, and very frequently mentioned in the free text replies. Access to support is mentioned less frequently, but the fact that access to the resources themselves occurs so often is a cause for some concern. This could be indicative of a lack of resources being available, though many of the institutions hosting respondents already have their own dedicated HPC centres, and will undoubtedly have access to both smaller local accelerators such as GPUs, as well as larger cloud compute resources such as AWS or Oracle Cloud.

From this we can conclude that despite the respondents repeatedly insisting that there main problem is a simple lack of resources, really the problem is that they simply don't know what resources there are. For example, Fig. 10 shows that many respondents require relatively simple features such as GPUs or multithreading. With the advent of cloud compute services it is trivial to gain access to such resources, at-least the hardware. There may however be problems obtaining system with specific level of security, design, governance or cost. Therefore, our conclusion here must be that respondents cannot locate sufficient support or guidance in order to either access or utilise existing HPC resources.

9.4 What is the perceived communication, accessibility and quality of the current HPC competence development, and do these develop further possibilities?

In general, there has been a positive response on the perceived quality. respondents tend to prefer in-persons event which makes accessibility more complicated. Locating courses and securing funding to

attend training are in general a moderate challenge. To have a more balanced offerings, there should be increased focus on practical aspects compared to theoretical aspects, and shorter+dense courses. As an additional note, there is an interest in self-paced learning that could be a stronger action point and more widely encouraged—including in WP8.2.

However, when we consider the partial conclusions from sections 7.4 and 9.3 we can say that although the communication, accessibility and quality are perceived as good, they have a major gap. This is in showing respondents what resources are actually available and how they can be accessed. It is hypothesised that this is due to HPC is a large body of knowledge, which will not be completely transferred to non-HPC specialists when they just need a quick and easy solution to their specific processing problem. Therefore, non-HPC specialists lack the necessary theoretical background to generalise approaches to marginally different resources.

9.5 Are there significant gaps in the current offerings that WT 8.2 can address?

The survey seems to indicate a wide interest in e.g. new features, but no single one dominating in numbers. Challenges are similarly widespread, with the most dominating issue being accessing compute resources. When considering the current course offerings, the perceived quality is overall rated above average, and the analysis did not uncover any significant gaps for the competence development offerings. If any, the survey indicates a single significant gap: accessing compute resources.

From this we can conclude that the main gap to address is in this lack of accessible resources, especially as this should not be a problem in the modern computing landscape. From these results it is unclear if this is a genuine lack of resources, or a lack of knowledge about how to access what resources are currently available, or a combination of the two.

9.6 Concluding remarks

This report into the survey of Danish PDC provision and training has presented a number of conclusions. The key findings were first presented in Section 1.1, with it hoped that it is now apparent were each has been drawn from. As a summary, this report identifies apparent shortcomings in existing material and training in allowing users to identify and access computer resources. Therefore we conclude that material and training needs to be added/updated to properly accommodate this. Training in HPC should also be presented primarily in shorter, denser courses with this being the format that most users expect. In order the generalize these results further, additional work will need to be undertaken.

Appendix A Question 5 and 6: Daily and largest compute resource

With questions 5 and 6, the respondents are asked about their current usage and experience with systems at various size with the questions "What computing resources do you use daily? Select all that apply" and "What is the largest system you have used for processing data?". The break-down is shown in Fig. 22 and 23

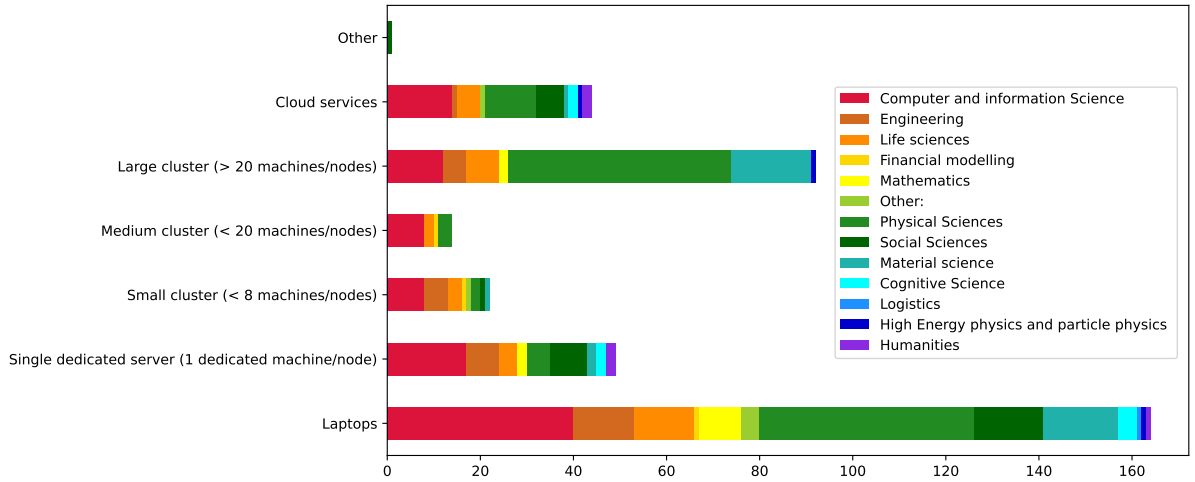

Figure 22: Distribution for question 5: Daily compute resources.

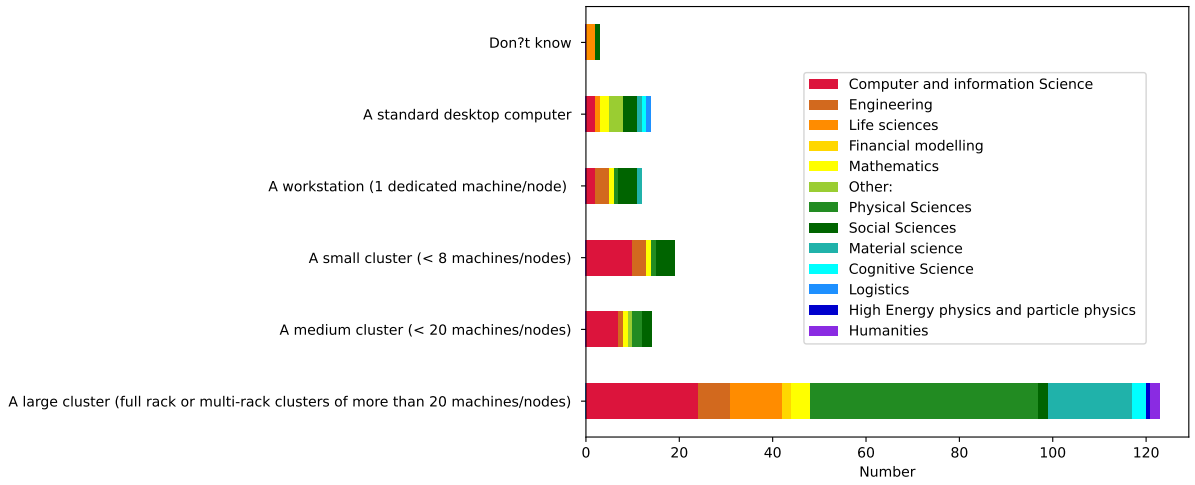

Figure 23: Distribution for question 6: Largest compute resources.

Appendix B Question 11: Features in your application

Question 11 addresses current features "What features do you use in your application? Select all that apply (leave empty if you don't know)" shown in Fig. 24.

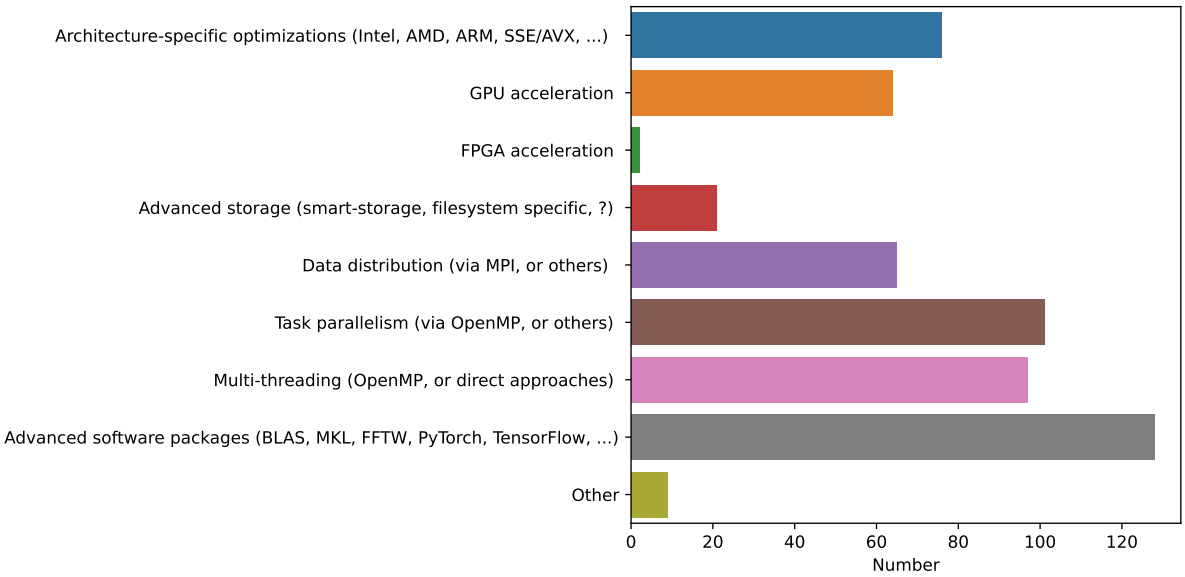

Figure 24: Distribution for question 11: Used features in application.

Appendix C Question 13: full page figure

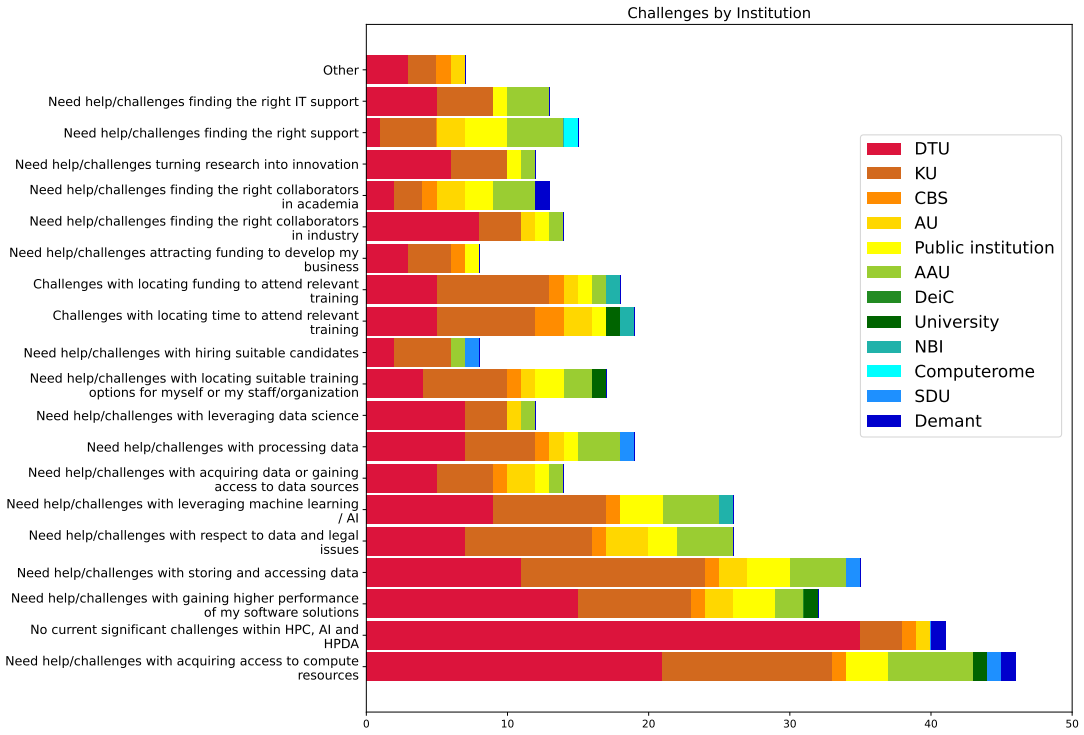

Figure 25: Response for question 13: Needs and challenges, broken down according to Institution.

Appendix D Question 13: Free text replies

NBI	Most problems boil down to lack of resources/ especially time
DTU	As experienced researcher in the field of of computational chemistry, I just need access to more powerful HPC Clusters, with high turnaround in the calculations and fast CPUs.
DTU	Support for live capturing, data storage and access of large amount of data in a specialized format. Ability to expose large amount of data in web applications.
KU	We use Stata and R. But I would highly benefit from another type of analysis such as web scraping. I also want to learn fundamentals of AI.
Public institution	Need time or stable manpower ressources to ensure systematic and proper data curation of large amounts of experimental data inc meta data.
SDU	I have my own compute and storage servers for data processing. For data production I have used SDU HorseShoe, Abacus, and now JoliotCurie via PRACE. (Type II HPC) There is very little support for personal HPC needs at SDU in relation to the SDU IT infrastructure. To share data with my group, I need the storage serve to be integrated with SDUs authentification service, support for such specialized issues is VERY slow, and can some times not be solved due to IT policies. An experimentalist can order support from the company supplying the equipment to help with technical problems. The same is not possible with HPC unless we create each our own world independent of the university IT solutions.
DTU	we are severely limited by the actual raw compute node available and not anything else
DTU	My #1 challenge is to find funding for development of open-source scientific software, and to make my expertise in open-source development part of my academic career. Neither my institution nor the danish minister for research has -to my knowledge- a plan for supporting open-source development for the benefit of the society and the local economy.
KU	ERDA is key for storing our large datasets. The local cluster at UCPH is key for long term research project where we reuse already obtained data. Funding is a challenge to maintain our local cluster, and national resources are currently not adequate. They need to be expanded by x10 to replace local resources.
AAU	Various constraints bind us to the Windows platform, and this limits the opportunity to use typical HPC facilities that are available.

Appendix E Question 22: Free text replies

In relation to your training needs and/or training experience (for example from your teaching experience), please feel free to write suggestions, ideas, criticism etc. that could contribute to strengthening the knowledge and use of HPC, AI, and HPDA in Denmark

AU	Step learning curve for linux operating system
AAU	No comments
DTU	Hands on workshops (online & offline)
DTU	encourage/enforce the use of linux at universities
Public institution	I have a large need that my students, that be master, Ph.D., postdocs will be giving an introduction to HPC, programing, and parallelization.
KU	We NEED Computerome 2.0 courses aimed at life scientists with little computational experience.
DTU	More hands-on sessions with experts, more theoretical background
DTU	I believe that common "code camps" or "hackathons" e.g. like those organised by https://www.gpuhackathons.org would be great to have locally in DK.
SDU	Training is not that important for me. I prefer to experiment with things myself in my own pace.
DTU	There is no need for additional training as the courses currently available at universities are sufficient.
DTU	we have an existing very high level of expertise in HPC
DTU	HPC training is often given by computer scientists with insufficient knowledge of the problem actually being solved, leading to solutions that do not apply to actual scientific programs.
DTU	Formulate a strategy for open-source development as a learning platform, as part of the available software infrastructure, and as a platform for entrepreneurship.
DTU	Create tutorials about what you can do with the software
DTU	Courses offered by universities provide a sufficient background
DTU	It would be great to have more courses on programming and HPC targeted researchers. Should be offered by the universites as Msc/PhD courses.
KU	The number and coverage of HPC courses in Denmark is quite sparse, and no central overview exists. It would be great if the DK node of EuroCC could help to establish that.
DTU	targeted training on specific projects is lacking
DTU	Courses offered by universities give a sufficient background.
DTU	Please notice, the priority rule should be modified based on different project.
KU	Especialyl wrt sensitive data, there's a lot of confusion. While e.g. SIFT should eb available to store and process sensitive data, the setup process and interaction is way too cumbersome. Now I rather apply for funding to pay for secure AWS or Google cloud storage where the acess is user friendly and easy and also the legal issues are solved since e.g. the S3 buckets are certified for secure data storage.
DTU	Currently, there are very few computational resources for undergraduate/graduate classes.
NaN	None
DTU	Courses offered by universities and interactions with colleagues provide sufficient background.
DTU	Local interactions with co-workers and maybe the occasional university course seems to work very well.

Appendix F Question 23: Free text replies

In relation to use of HPC, AI, and HPDA usage, please write any additional information about your current needs and what you perceive to be the current state of offers in Denmark

DTU	I am mostly working with COMSOL simulating systems with a large number of mesh elements. It would be good to know, how to more efficiently use the HPC and to optimize the performance for large systems, so they can benefit most from HPC and paralellization
CBS	I am besides working at CBS are associated with 2021.ai and I use their system also in research with gives me access to computation capacity but also ensures governance which is important in social sciences.
DTU	More CPUs are needed for running computer simulations of chemistry.
AAU	No comments
DTU	Competition for computing time and nodes is high; more raw resources will improve the rate at which we get results
DTU	We need more computational resources, especially for ML/AI. These systems should be set up carefully and appropriately so that students/researchers have convenient year-round access.
KU	Unclear what resources are available in Denmark and elsewhere, including Norway and Finland. Getting access involves too much bureaucracy
DTU	Massive CPU needs.....
KU	Considering how easy is accessing services like Google Colab, which allows the use of a GPU from Jupyter Notebooks, it is significantly more complicated to use GPUs or similar resources in Denmark (at KU) at the current stage. I recently started using UCloud. It seems to be useful, but it is a bit tricky to be thinking on the money that we are spending each time that we run a program. I think that as university workers, we should get access to computing power (including GPUs) and we should not be thinking about how much money we have. It is really cumbersome to be applying for a project at UCloud, or for extra time, each time that we want to run an experiment. Besides, it is really hard to know which are the resources that each department is allowed to use, as sometimes there are servers or HPC machines that are only for specific topics/department/faculties.
DTU	Needs of high VRAM (large models), currently very limited access to computational resources that can handle such models
DTU	Need more and faster CPUs.
DTU	more funding for hpc
DTU	Cloud based services platform-as-a-service
DTU	large amounts of CPU power
DTU	i would like access to a much larger HPC.
DTU	The all-overshadowing challenge is getting access to computational resources such as GPU accellerated compute clusters. The availability of such systems at large scale is poor.
DTU	more CPU needed
SDU	For my research, I need continuous access to DEIC type II HPC facilities. My work pattern is to run many simulations, but each simulation does not need to utilize a large number of cores. I usually use 192 cores per job, but are running jobs for a very long time (weeks), or many jobs for a short time. I use LAMMPS and have benchmarked GPU acceleration, but so far for my simulation models (Kremer-Grest for polymer materials physics), GPU does not offer any significant acceleration especially when taking the increased price of node hours into account.
DTU	We need much, much larger HPC resources to conduct serious machine learning research in Denmark.
DTU	We need CPU's not GPUs
DTU	we are severely limited by the amount of hardware available. Expertise is not being fully utilized due to hardware resource limitations. Denmark is as a whole also is far behind in compute resources available to researchers.
DTU	Clusters of Linux machines maintained locally at the university have historically been providing much more useful computational resources than centralized computer centers, as the funding is used to machines that are relevant for the problems of the scientists.
DTU	My needs software- and hardware-wise are currently covered. I rather miss a national strategy.
DTU	more computational power
DTU	Masssive amounts of CPU (not GPU) power are needed
DTU	Need more local computational ressources, that fit specialized local needs.
KU	The idea of having Type 1 "large virtual workstations" available for everybody in the academic landscape is great, but a global storage space is an important component for it to be a success. The amount of "traditional HPC" available through the national centers is highly inadequate, and support is missing. While I think it is good that national solutions exists, their current scale is too small, and does not cover the need of the community. Both in terms of performance, but also in terms of HDPA/AI. It is essential that the local HPC centers are not phased out, and large groups are allowed to maintain their own custom solutions (with own funding). The goals and roles are different and complementary to the national centers.
DTU	I need more computational resources
DTU	I require more CPUs
DTU	We need massive amounts of CPU power (not GPU).
DTU	Would require more CPUs.
DTU	high demanding of CPUs is required for study the realistic problem
Demant	My current needs are not large scale, long simulations, but rather short bursts where I do prototyping on different configurations or preprocessing of data. Short waiting times are therefore a priority for me. I am not aware of resources that are available to me, bar resources in my university, the company or renting "cloud compute".
DTU	There is no large supercomputing center in Denmark.
NaN	None
DTU	It's important that we have the flexibility to choose the right hardware and software and manage things locally.
KU	The majority of Machine Learning (AI) researchers just need access to GPUs that have easy access to their data. This is not really supported by DeiC. Therefore, each departments/university have their own GPU clusters in great numbers.

Appendix G Survey Lead-in text

EuroCC - Survey from the Danish National Competence Center

Dear participant,

This attached questionnaire is addressed to people that, like you, may use computing facilities and services.

Background
Computing, including but not limited to High Performance Computing (HPC), Artificial Intelligence (AI), and High Performance Data Analytics (HPDA), is playing an increasing role within a wide range of fields, from life sciences, physical sciences and mathematics, to medicine, linguistics and social science. The European Union has created the European High Performance Computing Joint Undertaking (in short EuroHPC JU), which is a legal and funding entity that aims at developing exascale supercomputers—based on European technology—for a wide range of European computing needs. Among the actions of EuroHPC JU, the EuroCC action, aims at building a European network of 33 HPC National Competence Centres (NCC). In addition, the Coordination and Support Action (CSA) CASTIEL promotes interaction and exchange between National Competence Centres (NCCs) in HPC-related topics addressed through the EuroCC project. Each of the 33 NCC will act locally to map available HPC competencies and identify existing knowledge gaps. The NCC will coordinate HPC expertise at national level, and ease access to European HPC opportunities for a wide variety of users from academia, public administration, and industry. DeiC represents the Danish EuroCC NCC (grant agreement #951732).
The work that DeiC is coordinating is divided into seven work tasks (WTs).

Reason for this survey
EuroCC WT 8.2 "Training and skills development" engages representatives from the University of Copenhagen (UCPH), the Technical University of Denmark (DTU), and Aalborg University (AAU). The long-term aim of the work of WT 8.2 is to facilitate effective training and skills development in the use of HPC, AI and HPDA in Denmark. The questionnaire has been developed with input from the EuroCC WT 8.3 "Technology Transfer & Business Development" and EuroCC WT 8.5 "Mapping of HPC/HPDA/AI Technical Competences" groups.

The survey is designed to identify which needs the Danish academic and industrial community has with regard to HPC, AI and HPDA resources and training. Therefore, we would be grateful if you could answer the questions listed in this questionnaire.

The data collected from this survey will be used in aggregated form within the EuroCC project and DeIC to evaluate and suggest training needs and competences.

Time needed for the questionnaire completion: About 10 minutes.

Deadline
We would be grateful if you could complete the questionnaire by September 15th 2021.
Privacy Policy
The collected data will be processed in compliance with the University of Copenhagen GDPR policy. If you specify your university in question 1, the results may be shared with your university.
Contact
If you have questions, please contact Allan Have Sørensen, alhs@adm.ku.dk

On behalf of the EuroCC team, thank you for your cooperation.

Appendix H Complete list of questions for the survey

1. Public institution or company (leave blank if you prefer not to say)

2. Number of employees in your organization
    - 0-9, 10-49, 50-249, 250-999, 1000+, Don't know
3. Your current job title (select the most suitable)
    - Associate/full professor / Senior researcher, PhD student / Postdoc / Assistant Professor / Researcher, Research software engineer / scientific programmer / developer, Managerial position in small/medium size company, Employee in large company, Employee in small/medium size company, Student / Undergraduate student, "If other, please specify"
4. Field of interest (select the most suitable)
    - Other:
5. What computing resources do you use daily?
    - Laptops, Single dedicated server (1 dedicated machine/node), Small cluster (< 8 machines/nodes), Medium cluster (< 20 machines/nodes), Large cluster (> 20 machines/nodes), Cloud services, Other
6. What is the largest system you have used for processing data?
7. Do you have in mind or do you work with project that would benefits from using more computing resources than you currently have available?,
    - a. If yes, do you have established plans for how you will develop your computing environment?
8. Do you need guidance or support for developing plans or explore options for your computing environment?
9. How many years of HPC, AI, and HPDA experience do you have? -
    - High Performance Computing (HPC), Machine Learning or Artificial Intelligence (AI / ML), High Performance Data Analysis (HPDA)
10. What best describes your daily usage and interaction with computer software?
11. What features do you use in your application? Select all that apply (leave empty if you don't know)
    - Architecture-specific optimizations (Intel, AMD, ARM, SSE/AVX, ...), GPU acceleration, FPGA acceleration, Advanced storage (smart-storage, filesystem specific, ?), Data distribution (via MPI, or others) , Task parallelism (via OpenMP, or others), Multi-threading (OpenMP, or direct approaches), Advanced software packages (BLAS, MKL, FFTW, PyTorch, TensorFlow, ...), Other
12. Which features would you like to add to your application? Select all that apply (leave empty if you don?t know)
    - Architecture-specific optimizations (Intel, AMD, ARM, SSE/AVX, ...), GPU acceleration, FPGA acceleration, Advanced storage (smart-storage, filesystem specific, ?), Data distribution (via MPI, or others), Task parallelism (via OpenMP, or others), Multi-threading (OpenMP, or direct approaches), Advanced software packages (BLAS, MKL, FFTW, PyTorch, TensorFlow, ...), Other, Need help with determining suitable features
13. What are your primary needs and challenges for yourself or organization? Select all that apply.
    - Need help/challenges with acquiring data or gaining access to data sources, Need help/challenges with acquiring access to compute resources, Need help/challenges with storing and accessing data, Need help/challenges with processing data, Need help/challenges with respect to data and legal issues, Need help/challenges with leveraging data science, Need help/challenges with leveraging machine learning / AI, Need help/challenges with gaining higher performance of my software solutions, Need help/challenges with locating suitable training options for myself or my staff/organization, Need help/challenges with hiring suitable candidates,

Challenges with locating time to attend relevant training, Challenges with locating funding to attend relevant training, Need help/challenges attracting funding to develop my business, Need help/challenges finding the right collaborators in industry, Need help/challenges finding the right collaborators in academia, Need help/challenges turning research into innovation, Need help/challenges finding the right support, Need help/challenges finding the right IT support, I find that I have sufficient competences and skills to conduct my current work within HPC, AI and HPDA (no significant challenges for the time being), Other, If you have additional information, missing opportunities, or can elaborate on the above, please clarify here:

14. Which of these support actions would you be interested in? Select all that applies.

    - Support with understanding what resources to use, Support with using advanced software packages, Support with selecting simulation software and other packages, Support with using computational resources, Support with using machine learning and AI, Support with using data analytics, Support with using workflows, Support with using High Performance Computing (HPC), Support with developing new software, Support with using accelerators, such as GPUs or FPGAs, Support with developing OpenMP code, Support with developing distributed MPI code, Support with performance debugging, Support with performance engineering, Support with using and accessing data, Support with implementing an IT system, Support with implementing and using a database, Support with implementing and using a storage system, Support with acquiring my own computational resources, Support with accessing petascale systems, Support with acessing pre-exascale systems, Support with developing my business plan, Support with commercializing an idea, Support with attracting funding, Support with writing funding proposals, Support with developing my business network, The support I receive is fully adequate, Other

15. In which of these areas, if any, do you provide support to your own and/or another organization? Select all that applies.

    - Support with understanding what resources to use, Support with using advanced software packages, Support with selecting simulation software and other packages, Support with using computational resources, Support with using machine learning and AI, Support with using data analytics, Support with using workflows, Support with using High Performance Computing (HPC), Support with developing new software, Support with using accelerators, such as GPUs or FPGAs, Support with developing OpenMP code, Support with developing distributed MPI code, Support with performance debugging, Support with performance engineering, Support with using and accessing data, Support with implementing an IT system, Support with implementing and using a database, Support with implementing and using a storage system, Support with acquiring my own computational resources, Support with accessing petascale systems, Support with acessing pre-exascale systems, Support with developing my business plan, Support with commercializing an idea, Support with attracting funding, Support with writing funding proposals, Support with developing my business network, Other

16. For which of these subjects, if any, would you be interested in training ? for you or your staff? Select all that apply.

    - Specific software packages for HPC, AI or HPDA, AI methods, HPDA methods, Linux operative system, including command line interfaces, scripting, editors, etc., Specific programming languages (Python, PERL, C++, C#, MATLAB, FORTRAN, SQL, etc.), Parallel programming methods (MPI, OpenMP for CPUs, OpenMP for GPUs), Computer simulation methods and tools (COMSOL, ROS, OpenFOAM etc.), Remote access and usage of HPC facilities (local, national and international), Software development tools and methods (version control systems, testing ...), How to write competitive proposals for computing time, I'm not interested in additional training. I would prefer to contact a support organization, hire or purchase expertise to do my data/computing work, Other subjects

17. Which of these training events, if any, do you find most useful? Select all that apply.

    - University semester style courses with classes over a longer time span (sparse), Workshop style that spans a few days (dense), Lecture types with presentations of material on board/slides,

Live demonstrations, Conference presentations, Networking events, Seminars, Guided exercises / hands-on sessions, Group work and discussions, Project work, Face-to-face tutorials by an expert (supervisor, scientific programmer, etc.), Self-learning (online-tutorials, books, articles etc.), Self-paced massive online open courses (MOOCs), Other

18. How many work-hours or work-days do you think is optimal for non-self-paced training events?
    - Work-hours, Work-days
19. Do you prefer in-person or online training events?
    - In-person, Online, Mixed in-person and online, Don't know, Non
20. Which types of training forms the basis of your current skill and competence level? Select all that apply.
    - Courses/workshops etc. delivered by universities, Courses/workshops etc. delivered by private/non-university organizations, Project work, Face-to-face interaction with an expert (supervisor, programmer, colleague etc.), Self-learning (online-tutorials, books, articles etc.), Self-paced massive online open courses (MOOCs), Other
21. If you previously have taken courses on programming languages, or courses regarding HPC, AI, and HPDA, please consider the last course that you attended.
    - a. Please rate your overall experience with that course:
    - b. Please rate the form and presentation:
    - c. To what degree did the course live up to its description:
    - Consider the same course you last attended. What issues did you encounter for this course. Select all that apply.
        - Too advanced in relation to my knowledge, Too basic in relation to what I need in my daily work, Too theoretical and not enough practical examples, Too practical, hence lack of theorical knowledge to develop further my expertise on the subject, Problems in applying the learned skills to my daily work after the course, Poor learning environment that hindered learning, I did not feel welcome to an extend that hindered learning, Significant technical challenges hindering learning, Too much material for the allocated time, My questions were not well-addressed or taken seriously, If any other challenges, please specify, I did not face any significant challenges,
22. In relation to your training needs and/or training experience (for example from your teaching experience), please feel free to write suggestions, ideas, criticism etc. that could contribute to strengthening the knowledge and use of HPC, AI, and HPDA in Denmark
23. In relation to use of HPC, AI, and HPDA usage, please write any additional information about your current needs and what you perceive to be the current state of offers in Denmark
24. We might like to get in touch for follow-up questions. Please provide your name and email address such that we can reach you for further comments. Leave blank if you do not want to be contacted.

References

[1] Sushil K Prasat et. al. NSF/IEEE-TCPP curriculum initiative on parallel and distributed computing - core topics for undergraduates (version 2.0-beta). `https://tcpp.cs.gsu.edu/curriculum/?q=system/files/TCPP%20PDC%20Curriculum%20V2.0beta-Nov12.2020.pdf`, 2020.

[2] J. L. Hansen, B. H. Jakobsen, F. Lo Verso, E. Molinaro, L. Nondal, and C. Pica. WP4 "Collaboration with industry" - Initial report on national HPC, Big Data and AI for SMEs. EuroHPC Competence Center Denmark, Jun. 2021.

[3] K. Thor Wikfeldt. Report on the 2020 ENCCS training survey. `http://media.enccs.se/2021/01/report-trainingsurvey-2020.pdf`, 2021.