Interview with a research associate and a principal investigator from swisstopo (Swiss Federal Office of Topography) at the Mont Terri underground rock laboratory (July 11, 2025)
Length: 1:30h

PM:

First, I would like to discuss individual actions with you again, what you can do in the application. I'll start with positioning. I'll leave the video running in the background and you can tell me what you thought of the functionalities and how you found it to use.

Expert 1 [Research Associate]:

So, I can comment on a few things while the video is playing. First, we had to select the area we were in and then the positioning, if I remember correctly. Those were the two spheres that needed to be placed. We know the laboratory very well. We can find these points quickly. I can also imagine that if these are really smaller niches, then others will be able to find them too, because in principle it's always done per scene (or use case/location). I can imagine that if the scenes were larger, or if you didn't need this (subdivision into) scenes at all, but could simply place it where you are, then it would of course be difficult to find these points. But otherwise, if you break it down like this: we had these two or three scenes, then I think it was relatively intuitive for me.

Expert 2 [Principal Investigator]:

It was very easy because you could tell exactly where it was based on the image I had in my left hand. If you have to orient yourself because you don't know where it is, it might take 30 seconds longer, but for me it was great.

Expert 1:

Yes, it was really smooth. We also had the option of re-referencing using all zone targets. Then it was actually relatively clean. Only in multiplayer mode was it different for the two users, or rather for the second user. The georeferencing didn't work so well there. There was always an offset of 10, 20 cm or so, as I remember. But in principle, it worked well.

PM:

I'll show you the video with this "refresh tracking." You have over 500 zone targets distributed throughout the lab. And when you notice that you have a deviation of maybe 10 or 20 cm here at the lamp [in the video], the idea was to reposition yourself in passing using two zone targets. Your infrastructure is actually great for this mechanism, as there is a zone target every 5-6 m. However, you can also see that as soon as you reposition yourself, the last ball placed is again 5-10 cm "off." There are three reasons for this: First, there is a drift as you move from one point to another. Second, manual referencing is only accurate to a few cm. When you position it by hand, you can only do so to an accuracy of a few cm. And third, unlike the fixed points in " ," the zone targets are not 100% accurately calibrated. But we could discuss what level of accuracy is actually required.

Expert 1:

We were limited by what is permitted in this interface. These stickers are also coded, of course. The DataMatrix code is coded to this zone target number. It would be nicer, of course, if it were a kind of automated process. So that as soon as I look at one with the glasses, the code is recognized and the coordinates are set. We have now set these spheres manually. Perhaps there is also a comparison that says, "Yes, you have to set the coordinates here on the plate." Of course, the plate also has a certain size, but if the process is even more integrated, so to speak, that I don't have to do it manually, then I can imagine that there would also be opportunities to optimize it. You probably have to get relatively close to such a zone target with these glasses, but then this process would of course be even faster. But I don't think this option is available [in the current application] with these DataMatrix codes or this type of recognition. But the spheres worked well. Especially for the first user, who was able to do it relatively cleanly.

Expert 2:

I was actually extremely amazed at how stable it is when you move.

Expert 1:

We also walked through the lab. The shift wasn't several meters, so you wouldn't know where anything was anymore. That was good.

Expert 2:

Imagine hanging the entire coordinate system at a height of 5 m, at two points, and then moving 300 m and ending up 30 or 50 cm off. That really surprised me. I hadn't expected that.

PM:

I think, [Expert 1], that your suggestion is definitely feasible. The application can recognize QR codes directly. That's the use case we'll discuss later, where the glasses automatically open the Fulcrum window [productive database] when they detect a QR code. I think you can do that with positioning as well. I'll open the video with the QR codes... The camera on the glasses checks where there are readable codes. According to the library used [xzing], these could also be DataMatrix codes. I was a little surprised why these are not reliably recognized. In any case, it works most smoothly with QR codes. After recognition, a button is displayed that you can click on. The button is exactly in the middle of the QR code. That's why I believe that this positioning can also be automated at the beginning using this principle. Of course, you have to get a little closer, maybe to a distance of 30 or 40 cm.

Expert 1:

That would mean that when I go to the code that could be detected, a button may or may not appear. And then I could somehow have the option to click on it so that it uses the code for referencing or for displaying information. So you could have different options. It's good for us in the lab because these zone targets are always at eye level. You can get relatively close to all of them. And if you can do it more smoothly and don't have to position the spheres manually, you save a certain amount of time. You would do this again and again if you were traveling longer distances [with the glasses].

PM:

What did you think of the feature that automatically opens the Fulcrum entry as soon as you detect a QR code on an object?

Expert 1:

When we do a follow-up [subsequent integration or updating of BIM models] or objects have different statuses and you want to edit them further, it's obviously a useful feature. So you can go into Fulcrum and edit the entry. We didn't do the editing ourselves, and I think it would be relatively time-consuming because you would have to set up these windows for an object and then scroll and look. You can do that for individual objects, but I think it would be too complicated for everything. That's because we have access to the backend. We don't have access to the interface of the mobile applications. Then, of course, it would be a process that we would normally do.

Expert 2:

With this "gun keyboard" [allusion to selecting individual letters with the controller on a virtual keyboard], you don't want to write anything at all...

Expert 1:

I'm more concerned with selection fields: changing a status, yes/no options. I think things like that work well directly. But with text, you're naturally much slower.

Expert 2:

I have the impression that if you want to change a Fulcrum entry and have to write something, you can do that easily with your smartphone, even while wearing the glasses. I tried it. You can see well enough, which surprised me.

Expert 1:

Yes, that would mean that we can recognize where to work in the XR space. You can coordinate spatially well. And then you can use the tablet with the Fulcrum app again. One workflow we definitely need to consider is damage mapping of objects. It would be nice if we didn't have to use the Fulcrum interface because it's complicated to navigate with the glasses. But if you do mapping for simple cases that can be displayed in AR, that would of course be very good. Then you can do it much faster.

PM:

We can also jump straight to measuring and drawing [the video is called up]. Here, you could switch between different modes. There is a mode where you saw a blue preview sphere. And then the camera used the depth sensor to see where there were points of contact between this invisible beam emanating from the controller and the real world. Click once to place a point. Click twice to place a line. Click three times to place a triangle with the area and the lengths of all outer lines, and click four times to place a quadrangle. The points can be placed reasonably reliably at distances of perhaps 6 to a maximum of 7 meters. You can measure all kinds of surfaces—whether on the ceiling, walls, or floor. And then there's this second mode, where you have a red preview sphere. This allows me to spray graffiti virtually.

Expert 2:

Are these graffiti located as 3D objects? Or is it just paint or pixels?

PM:

They are actually located, yes. They exist [in Unity] as so-called line renderer objects. So ultimately, they are lines with lots of intermediate points. And you can simply save and export them as 3D objects in perspective. That means you can get the drawings out of the application again.

Expert 2:

You then get polylines, so to speak.

PM:

Exactly, a polyline.

Expert 1:

I can imagine that once you're at the stage where you're using the glasses to complete applications and workflows, you might want to have these functions available in the room as well. I can measure something, I can mark something specific. When we're working on site and need to measure something, we're not going to put on the glasses just to measure two or three meters. But once you're really immersed in the XR environment and doing a lot of work there, measuring is obviously not a bad option.

I can measure the room: Is there enough space for a new cabinet? We can pre-mark where something should be placed. In the past, you would just walk in and mark or spray something on the floor. But then you couldn't get it off. That's the advantage here [with the XR application]. But of course we work with many other groups. That can be difficult at times, because everyone would have to work in the virtual space, so to speak. But the functions are very useful. I haven't tested how good they are. We didn't lay out the measuring tape and remeasure. We didn't do that. But I find these two functions very useful.

Expert 2:

You probably measured this grate on the floor, this manhole cover, right?

PM:

It was exactly 10 m^2. I played through the scenario shown in the video twice. Once it was – I think – 10.07 m^2 and the other time 9.95 m^2 or so.

Expert 2:

So you checked that. Then that's fine, great.

PM:

I think it will be pretty accurate to within a few cm.

Expert 2:

That's enough, because a spray like that is only accurate to a few centimeters anyway. It's certainly useful for getting an idea of how much space there is or how it works. However, I could also do that in a model where I don't have XR. It's definitely not useful for me to discuss with the people who are doing the drilling.

Expert 1:

Another point comes to mind. Until now, our planned workflow when placing new objects has been to laser scan the location, then georeference it, extract the coordinates of the objects from this laser scan, and then create the object. I think this is relatively precise in terms of coordinates and positioning. You could think differently if we already had the geometries. For example, we know the length, width, and height of a new cabinet. We make it available as a virtual object and can then place it in the room. Either it is already there and we adjust it visually, or we simply place it on the 3D virtual coordinates. This could be a faster workflow because we can already tell from the companies what the coordinates and dimensions of the objects are. Then we get them, create these 3D files, and place them in the virtual space. Then we don't have to take the detour of doing new laser scans and so on. However, we have a certain tolerance here, so there will probably be deviations in the cm range. But that could, of course, greatly speed up this workflow because we don't have to scan and extract coordinates outside the lab and then place the object. I could also imagine us doing a workflow like that. It's similar to what you just showed with the measuring—except that I simply place an object in the space and position it.

Expert 2:

I think it's a great idea, [Expert 1].

Expert 1:

It's just super fast. We get a list of new objects and then we go there and place them. You could somehow program functions so that it has to stand on the floor and doesn't float in the air. Or so that a plane is defined where it is placed. On the wall, for example. So we can think more about what we need. But I think that would be a very fast workflow. And you could also use it to plan experiment rooms, of course. You can see: Where does something fit? How much distance do I have from this cabinet to the wall? I would find these functions relatively interesting for our processes. The thing is, we can only scan once the objects have been placed. Maybe one will be placed today and the next one in a month. Then, of course, you wait until they are all there. And [with the XR application], you could work relatively quickly.

Expert 2:

We simply run into the problem that we are too inaccurate with this application in terms of coordinates. If you want to put it into the BIM later, you need the correct coordinates. Otherwise, you have this ten-centimeter offset. We have our zone targets, whose coordinates are not accurate. Because of them, we are not precisely located. The dimensions themselves are correct, because it doesn't matter for the proportions whether there is a certain shift in them. That is not distorted. But if we take coordinates from them and put them in the BIM, I can imagine that we will be off there.

PM:

There will definitely be a certain deviation. We would have to see how big it is.

Expert 1:

We would have to look into that. In principle, that would be for the update. That's a valid point, but I don't know enough about how high the tolerances are that we want to accept...

Expert 2:

In BIM with the installations, I would of course like to be accurate to within 1-2 cm. And with this shift that we now have in there, it looks to me like it's more than 2 cm. But I'm happy to be proven wrong.

Expert 1:

We would have to see whether this georeferencing can be further improved. And maybe at some point we'll reach the point where we can say, "Yes, this tolerance is OK. It's somehow less than 2-3 cm." That could well be the case.

PM:

I think two factors could lead to smaller deviations in the future. First of all, if you implement the glasses to automatically recognize the code and then find the exact center point. This will probably be more accurate than manually positioning the red balls. And the other point is that the glasses receive regular software updates. From the moment you position yourself, the built-in hardware, such as the IMU, ensures continuous tracking. From the point where it starts, it's a kind of continuous error propagation. Presumably, the camera data is also used for this purpose. I was quite surprised, because I had already done similar experiments in Mont Terri in December and now noticed – probably due to software updates – that the accuracy was better. So this is an area where the manufacturer still has some room for improvement. If there is ever a Meta Quest 4 or something like that, I assume that the tracking will be even more accurate. Whether you end up with only a 1-2 cm deviation, I can only speculate at this point. But I believe that in the future it will definitely be better than the accuracy we have achieved here now.

Let's now talk about the presentation of the semantic data, the "tree structure." You, [Expert 2], were the statistician for a video that was attached to one of the drill holes here. I'll start the video. Here you can see how to navigate through this data. Essentially, it was your suggestion, [Expert 1], to display the data by category. How do you think this was implemented?

Expert 1:

When we select and open individual objects, we have very different information. So, classifying or categorizing them accordingly with these disks [meaning the discs that represent categories within the tree structure] and then listing the topics below... That was well implemented, it was easy to read. I think that when you then accessed individual pieces of information, such as the image data, each photo was listed individually. For example, if we have samples, each sample was listed individually. And it would be good to continue working with lists, because we might have 100 samples and then you don't know where exactly everything is. For the overview, you would have to work differently in the next level. But this structure, where you first list the category and then go into detail, makes a lot of sense. And I think it was also very well implemented. You could quickly grasp what information was available. If one of the disks wasn't there, then that information wasn't there either. Of course, you could always display all categories for an object class, but then color-coded—e.g., the blue ones have data and the white ones don't. That way, you always know that there really is no data and that it's not some kind of programming error. But overall, it's been solved in a very sensible way. The samples are different objects from the drill holes. And you combined them based on the drill hole number. It worked well for me. The little arrows underneath, which you use to scroll through the list when data is available, were a bit small.

PM:

That's right, yes.

Expert 1:

You had to try for a while before you could move on. Accordingly, it might be good to have a list that opens up. And practically a kind of scroll bar. Then you can find things a little faster. But it could also be like a kind of photo board, so that you don't have to click on the photo first to see it, but can see all the photos immediately in a thumbnail view.

PM:

Yes, a kind of gallery...

Expert 1:

Exactly. With samples, it's often about how deep they come from. And then, if I want to see the photo in large format, I click on it again. Or I can mark the three or four things I want in the first step. Then I open them all at once and the view is already coordinated in a way. Because now it was like this: every time I opened a new panel, it was new in the room again. And then it takes a relatively long time to position them individually so that you can view them. So I see a lot of potential for development there, a lot of room for improvement. But the strategy of dividing the data into topics is, of course, very good.

PM:

I based the categories on Fulcrum. There are also these higher-level categories there. I had summarized some of them, but basically I found it relatively easy to see what suitable categories might be for each object. And yes, in terms of user guidance, as you say, there is still room for improvement in terms of intelligence.

Expert 1:

One thing I noticed: a mouse-over function would be nice. When you enter a workflow, you might want to know something immediately about drilling operations with the help of a mouse-over. And not only after clicking on the categories and so on, but I want to know immediately when drilling took place, how it was done, and so on, so that the initial information that is defined is displayed immediately. That way, I don't have to keep clicking individual windows on and off, but instead I can use a mouse-over to read most of the information, and if I want to delve deeper, I can move on to the next step, which is to open up the tree with the individual topics, so to speak. Of course, this is just an initial test. If we start from a specific use case, we would probably try not to work with so many panels.

PM:

Exactly, then you can tailor it to this use case. And yes, that's a nice idea with this preview. It saves a lot of clicking.

Expert 1:

The information you want to see is probably very similar for many objects, but for some it is more specific. And what I don't think we've seen so far is that all the BIM objects have a status indicating whether we've edited them or not. And it would be nice to implement something like that, of course, so that you can see the status. That can be done with colors alone, using the virtual objects. Or filtering by status—those are the kinds of things we can work with well.

Another thing that occurred to me: If we say that there are fields for some objects that absolutely must be filled in but haven't been yet. Then information is missing, and that's exactly what my filter is for. I don't want to say, "Oh, now filter this and that for me." It's very time-consuming to create a clever filter that shows me all the fields where there is no information yet and which need to be filled in. That would be so much faster, because then I wouldn't have to [navigate through the entire tree structure] first; instead, the filter would show me all the fields that I still need to describe.

Then I click on the object and receive the information that, for example, three fields are still open here and I should fill them in. Something like that would of course be much faster than filtering again for each new object class. If we go in the direction of continuing to update the metadata in the lab with some kind of clever interface that maps back to Fulcrum, to this main database... So still allow free filtering, but for this updating, the filter should perhaps be set automatically to the information that is still missing.

PM:

Yes, that's a nice idea. If you have a specific use case in mind, you can tailor the application and user guidance precisely to that. Here, the approach was more that we have a lot of data and don't yet know exactly what we want to do with it. The first step was to get a basic overview and provide basic functions for this.

Expert 1:

In any case, it was very good for us to see what options are available. [In the next step], however, you have to implement it in such a way that you can really say, "Hey, this is really cool, I'm faster." Because in the end, it's all about time for us. And if you have to filter for a long time, search for a long time, or the buttons are too small, then you get annoyed because you can't move forward quickly enough. You can do a lot, but you're just not fast enough.

PM:

[Expert 2], do you have any thoughts on the tree structure?

Expert 2:

Yes, well, at first glance, I found it a bit too complicated. But of course, that's something you can figure out as you use it. Of course, every time I clicked, I had no idea what would happen next and how to proceed. It's also a little different every time, because sometimes a lot of things come up, sometimes only a few. I just don't know what will come up for whom. Then there are [the categories] "General" and "Core," and then it gets blurry again and is almost impossible to read. So now only here at "Core," otherwise it's fine. For me, everything was also a little too small visually, it was difficult to read. But it was actually performant—when I clicked on it, something happened. I just never knew what would happen. That was the first time and you get used to it. My use case, which I saw right away, is much more basic than what [Expert 1] discussed earlier. For me, it was just extremely helpful to see all the BIM elements in 3D to see what was missing and what was wrong. Are the drill holes there? Are the assets there? Are they labeled correctly? So if these QR codes work in such a way that I can see what the BIM says about this element, that's very useful. Of course, this is more useful at the beginning of our BIM, until everything has been reconciled. At some point, it's no longer an issue. Nevertheless, you always have to check whether it's correct. So I do see a certain benefit. But I noticed that After about 20 minutes with the glasses on my head, I didn't find it so great anymore. Not because I felt dizzy or anything, but because it was very tiring for my eyes. Everything flickers a little, everything moves a little and then comes back into focus. It was very tiring for my eyes. Maybe I'm just too old, but I found it really exhausting.

Expert 1:

It was fine for me. But I can imagine that I wouldn't want to do something like that all day. The other thing is the flickering. Especially when you're reading text or something that's very small. That can be very stressful.

Expert 2:

You just have to concentrate extremely hard when reading. Your eyes have to focus constantly.

Expert 1:

It's probably really tiring if you do it for a long time. There's one more thing I want to mention before I forget. There was always this intermediate window. We clicked on an element and then an intermediate window appeared with various options. But you actually want to open the element immediately. What else I could do in this intermediate window with the six or seven buttons wasn't very intuitive. That was an extra thing. I have to open it every time. I still have this second window every time. So which one belongs to which?

PM:

Yes, that's right. When you move them apart, you can no longer see the connection.

Expert 1:

So I would somehow delete that from this hierarchy so that it's implemented differently. I found it relatively complicated when I click on the photo and it doesn't open right away, but there's an intermediate layer and I have to open it a second time. And you have so many buttons on the controllers. At first, of course, you don't know which button does what. But you could imagine programming them so that you can do one thing with one button and another thing with another button. For example, one button is for opening and the other is for viewing directly.

PM:

In this video clip, we can see the many photos of the samples. There is a frame around these photos. Theoretically, there would be enough space to accommodate all these buttons from the intermediate view. You could do it with an extra button or even here in the photo that is then opened.

Expert 1:

Or, instead of zooming in on the photo when you first view it, you could just show it in full size so that you can see it right away. I think you can optimize each individual element significantly. It works relatively quickly, which is what [Expert 2] said. You click on it and it's there immediately. There's no waiting time, which is perfect. But how it's displayed, in what size, in what arrangement, so to speak, when you open several elements... When I open one and then the next, I actually want the second one to be docked to the first box in the correct orientation. There are many options for improvement here. Or close all elements at once. If I have too many, I don't want to have to click on every single "X" to close them...

PM:

... Yes, that's right, it's not as much fun as placing them [laughs].

Expert 1:

There could be more user-friendliness there. But of course, that's something that would have to be specifically programmed or perhaps something that the glasses manufacturer will allow at some point.

PM:

As I said, you can basically do anything with the Unity game engine and the glasses. You would have to implement it yourself, but it is definitely possible. I call these "comfort functions" – there are no technical hurdles, you just have to sit down and do it.

Now we have a few more features to discuss. I'll start with searching and filtering. There is this time filtering feature. The motivation for this feature comes primarily from the decommissioning of nuclear technology facilities: for this process, I need to have access to the entire operating history and know what happened to which object during a specific period of time. So it's about displaying different levels of information. I aimed to achieve this principle here with the TIS [Tracer- e Information System]. The date is known for almost all injections from the TIS. In the video, you can now see how to filter all injections made before the year 2000. I can enter this using a virtual keyboard. The application then displays only those objects that already existed before this date, in addition to the injections. This was my original motivation for this use case here. It should be possible to easily trace back what the information status was x years ago. How useful did you find this filter option?

Expert 1:

I think we also need a lot of filtering options. With a time filter like this, i.e., to see which objects existed back then and which exist today, entering the exact date may be complicated. You could consider a time slider that you can move to the right and left to show and hide objects. Or you could just look at what was added in the last year. Of course, that works just as well. So you would need a time range.

But then it always comes down to how I can do something faster. And typing in a date on the keyboard is perhaps much slower than if I have a time slider where I first slide the year, then the month, and then the day. The same applies if we only want to see specific experiments. Here, too, you could type in the abbreviations using the keyboard or have a selection tile or something else. The more filter options there are, the better. But you have to make them so that you can filter quickly.

Expert 2:

I definitely think it makes sense to filter by experiments and view everything that goes with them. We've already talked about that. Whether I then have to click on those three letters or have a selection list doesn't really matter to me. There are so many experiments that if I have to slide up and down, I might be just as fast as if I simply typed in two letters. I also think it would be good if you could filter the asset types so that I only see "power" or "network," for example.

Expert 1:

With so many menus, I wonder if the manufacturer now allows audio filtering. In other words, AI that says, "Show me only the CI experiment." And then it does that.

PM:

Yes, that's possible.

Expert 1:

I think things like that are extremely important, and of course they have to work reliably, but they would speed up the whole process. I don't need all these input masks, and [the glasses] have a microphone. I would also be very interested to know whether filtering can be done with AI. Then I can quickly show or hide what I need. Before looking closely at the use cases at our company, we would have to take another look at what can actually be implemented technically and what the options are. Regarding the slider—we discussed this in relation to this one field, I believe in the CI experiment, where all these holes go down into the ground. I was there with [a colleague] and we said to each other that this would be something for a visitor group at . It would give them a better idea of where each borehole was drilled, when it was injected, when it was overdrilled and sampled, or whatever. With a time slider like this, I can of course also easily follow the progress of an experiment, because I can see what came before and after. This is an application that has already been seen in many such applications. I also think that swisstopo in the GeoDataLab in Bern has a time slider function to see what the glacier level or the history of land use was. Time sliders are of course very clever for such applications.

PM:

Basically, it's all doable. The glasses have a microphone, and you can also use gesture control. It actually works. If you put the controllers away, you can use your thumb and index finger to simulate a kind of cursor. But when I was testing it, I noticed that it's just much more comfortable and precise with the controllers. That's why I didn't look into gesture control any further. In some projects, we did a lot with Speech recognition, and it definitely works. You just have to define what the user can say and pick up keywords. The application has to know that when I say this keyword, this particular function is associated with it. So you have to do a kind of mapping. That's basically possible.

Expert 2:

But that can only go wrong...

PM:

It works with a large language model [LLM]...

Expert 2:

... With an LLM, you don't have to think in advance about what it might say. So if you now have to define keywords that I could say so that a function is called up, then I have to say exactly what you have thought of. That can't go well.

PM:

But that's actually how you have to do it. An LLM only returns text as a result and doesn't really know whether it's correct in the overall context. We just have to interpret it. And if an application is supposed to interpret an input from us, then we have to specify what's behind it. If I say, "Filter all drill holes that were already there before the year 2000," then the application has to be able to interpret that. It will call up this kind of timeline in the background; it's about the object class drill holes, and both are to be combined via filtering. So this mapping would have to be done. But I think it's actually quite simple because, unlike a real LLM, completely crazy inputs are not to be expected. We're not asking for the best recipe for tomato soup, but rather everything is limited to the use cases that can actually be done with the application. So it's not that far-fetched.

I'll open the video with the element search. You may remember when you searched for "Sample 25" and then all objects that have a Knowledge element with that name were listed. Currently, the application only searches the title fields and description fields, i.e., everything that is in this intermediate window. This could also be extended to linked documents such as PDFs or perhaps even images. That would then be RAG, or Retrieval- Augmented Generation. This would enable me to access a database in a similar way to an LLM – but not a public one from the internet, rather a database stored here. In a research project, colleagues of mine are working with RAG [details about the use case and project partner are provided] and developing a type of LLM whose knowledge base consists exclusively of stored documents. I can well imagine something like this here, because it is clear which object classes exist and which objects are behind them.

Now we've come to the last use case with navigation, i.e., being able to select objects and navigate to them. How intuitive was that for you?

Expert 2:

It was great for me. The only annoying thing was that you get kicked out every 20 meters [Background: The Meta Quest 3 displays a warning message that must be confirmed with OK when you have moved far away from your starting position. Only then can you see the application again]. But the guidance, i.e., the navigation, was good. In the last few meters, it didn't really guide me to the hole. But I was so close that I could already see it, so it wasn't a problem. It worked well.

Expert 1:

I don't think any explanation was really necessary. I didn't quite understand why the button is called "Linked Object" and not "Guiding" or something like that. Then the arrows were there, and after that I didn't really need any more explanation. It was also clear what the different colors of the arrows stood for. I thought that was good. You could also see the distance at the beginning. You said that this is the distance as the crow flies. Of course, it would also be nice to know the direct walking distance, or both pieces of information. Or while I'm walking, I can see how many meters I still have to walk. But in principle, it was simple. What I didn't like so much and what can be a bit irritating is that the arrows are visible in the next gallery, even though I can't see them visually in the room. Sure, I can see that now I have to turn right. That has an advantage, but somehow it was a bit irritating. And of course, the app makes the arrows show the shortest route, so not always in the middle of the gallery. I wouldn't walk that way, because I don't walk so close to the objects. These are just small details that don't have to be annoying, but maybe there are a few things that could be considered. Some people then walk exactly along this arrow. And then they walk too close to the equipment. Especially with visitors—if you tell them, "Oh, look, there's a scavenger hunt, now you can walk that way," they walk into objects. But this case, where I'm being guided somewhere, is basically brilliant. I think it's great. I can also imagine many things you could do with it.

Expert 2:

Personally, I like being able to see where to go next. In other words, I like being able to see the invisible path. I was more pleased than annoyed by this. I understand your point, [Expert 1], because if you're not familiar with the area, it can be confusing. If you are familiar with the area, it helps me decide whether or not to trust [the navigation].

Expert 1:

Sure, you can leave it as it is and give the things that are not in the field of visibility a different color. But you could also perhaps display a map that tells me where I am in our laboratory area. I actually thought it was great to know that the way is back there. At some point in the " " phase, I think we discussed whether this was really the shortest route and whether it would take us to the right place, because the perspective looked wrong. But it turned out to be correct after all. It only confused me briefly. So you misjudge the distances or what is behind and in front of you. But in itself, being guided is of course really great.

And just being able to see what objects we have and whether they are the right ones. This is of course very useful for someone who is new to the lab and doesn't know the equipment yet. They can walk through and immediately get information about what each piece of equipment is. They don't have to go up to each thing and scan it, they can see it right away. I think that's a really great application, just placing these objects and providing explanations. And of course, if a little more information were displayed above or next to each object, then we would actually already have a use case, namely getting to know the laboratory instruments or objects by walking through them. And this function can also be used to guide the person to very specific installations. I can well imagine that people would want to use something like this very quickly.

Expert 2:

Just the idea that [a colleague] is on-call service and I have to tell him where to go. I now do that with a FaceTime call, where I say to him: "See that up ahead, now you have to go right, now you have to go left, now behind it." And otherwise, he enters it into his glasses. If he can do that, then he'll be guided there. Yes, that's cool. I like the idea of the map, [Expert 1], because I'm a map lover. But I'm not sure if it will help the average user. People can't read maps.

Expert 1:

I have two more topics. First, we had to take our board with us somehow.

PM:

Yes, the "virtual backpack" [meaning that higher-level windows such as the search and filter window remain in one place unless they are taken along virtually].

Expert 1:

I'd like the backpack to be a bit smarter somehow. The other thing is: now I have the guide to some object, of course. And I'm not just going to one, I'm selecting 30 different things, so I want to be guided there in a clever, optimized way. Can it do that? That's one thing. I'd be curious to know. The other thing is, if I decide against it – "Oh, now I'm going to the other one first" – as in classic navigation: does it recalculate the route?

PM:

Yes, it could. This navigation system in Unity is quite flexible.

Expert 1:

But would there also be another route? So if I go through the other gallery and not through the predetermined one...?

PM:

It replans. It always calculates the shortest distance from where you are to your destination. And if you go somewhere else, it may navigate you through a different gallery.

Expert 1:

Because with a picket like this, when you say, "Here you have these 60 objects, you have to go to them, you have to do something there." Maybe you have to go somewhere else or something happens. You have to think through all these things, but I can already see interesting applications for getting to know the lab or for checks that you do.

PM:

Regular tours could also be stored, yes.

Expert 1:

And if there are two or three pieces of information that need to be provided for each object, i.e., actions such as "check," "uncheck," or a comment, and then that is mapped to another database... Then you can send out our report, [Expert 2]. Then you're a bit faster, of course [meaning the creation of the report on picket service intervention in the underground laboratory].

PM:

Here's a follow-up question, because we discussed earlier, based on the navigation, whether it's better to see invisible objects or not. That was a question I kept asking myself. Here, for example, in the current video clip, it's still okay, but you can also see that most of the objects that can be seen should actually be invisible. Conversely, especially with hidden objects such as drill holes, this is the only way to make them visible. But I have also often asked myself whether this should perhaps be limited to a certain distance. Perhaps you would want to see all objects within a radius of maybe 20 meters and initially hide the hidden objects outside this distance. This would prevent confusion with so many virtual objects far in the background. Could you please discuss whether you found this distracting?

Expert 2:

Well, I didn't find it distracting, and it didn't confuse me after two seconds. What was annoying were those blue balls that represented the injections. They were incredibly obtrusive, even at a distance of 200 meters. You could see them very clearly. I actually think it makes sense to see the other objects. Of course, you could consider a filter, but 20 meters is definitely not enough. You could also consider simply dimming objects at a distance so that objects far away are less bright.

PM:

With a transparency gradient, perhaps...

Expert 2:

... Exactly, so that a hole in the background is less visible. You can still see that something is there, but it doesn't catch your eye as much as what's right in front of you. Some form of fade-out wouldn't be a bad idea, I think. But I also think it's cool when I'm standing there in [a certain gallery], looking at the wall and seeing the holes at the very back. But then I also see [another gallery]. And for the first time, I saw how close the [certain holes] are to the [other holes] in the gallery. I think it's important to be able to see through the wall. But when I'm standing in one gallery, I don't always have to see [a certain hole] in the parking niche. It's too far away. But 20 m would be too little for me. I would rather have it in the range of 50 to 100 m. And work with a gradient, if that helps.

Expert 1:

I can well imagine that some scenes are overwhelming because, for example, large objects appear to be close by. But they are simply large and long and somewhere far away. Perhaps you could work with different options. We talked about the "virtual backpack" earlier. And you also showed us the photo of where we have to walk to when georeferencing. So if I had a menu on my left hand where I could configure things like show/don't show, thread up to a certain distance, show additional information, etc., that would be great. You would also have the option of showing or hiding more information about objects without having to walk directly to the object. And the same applies, of course, to taking some panels with you, where you say, "OK, please pack that for me, I want to keep it even though the object is far away." But anyway, we'll just have to look at the details to see what specific use cases there would be for Mont Terri. I also think that everyone has a different feeling about it. Some people are probably really irritated by the fact that things are displayed that aren't actually there. And I think the distances and sizes can also be very confusing for some people. It was OK for me. I would just filter more. When I have a task to do, I only want to see the things that I need to be interested in right now. Then you're simply less irritated. When you're in the planning phase, you want to generate more background information so that you don't put something somewhere where there might be something important, such as a sensor that you don't have in mind at the moment. But I think it was good so far. You also get more information and you might see what else is nearby. It's a mature lab, after all. There are also very old structures that might be there somehow, but you don't really have them in mind.

Expert 2:

We haven't mapped any sensors anywhere, have we?

Expert 1:

No, we don't have any sensors here. That would come later. Right now, we're just looking at these BIM objects and some geology in the drillings. Of course, it's interesting. We discussed where these great drillings can be seen. If we also had the Main fault [a tectonic fault zone that runs through the laboratory] as a 3D model, the context would of course still be there. So we can still think further.

Expert 2:

Well, we do have the geological areas.

Expert 1:

We do. But such areas can also be confusing. A borehole like this is less distracting. But if you have something flat at a great distance, I think it becomes relatively complicated to understand.

PM:

There was one more feature that we hadn't discussed yet. That was the access history. So, you could see which user had last viewed an object. Every real user had a kind of name tag above them in multiplayer mode, and you could click on it. Then you could see which objects and elements that person had last interacted with. Is this a useful feature?

Expert 2:

Well, it reliably crashed my game every time. To be honest, I can see the person standing next to me. I don't need a sign above their head. And I don't really care about their viewing history either. It's only interesting when I see the history of what was last changed on the individual objects. But this was only about what the person wearing the glasses had last looked at in the lab. I don't see why I would want that.

PM:

The motivation behind this is the concept of Knowledge maps. In large teams, it can be advantageous to create a Knowledge map about yourself. When people work remotely a lot, or even in larger companies in general, you often don't know who is in other departments and what knowledge they have. This is often synonymous with the topics they have been working on recently. For the application here, I admit that I took a very simplistic approach by depicting such a Knowledge map in tabular form. And yes, of course, you could just ask people, and there's certainly something to that. But my goal was simply to bring the concept of a Knowledge map into an XR environment.

Expert 2:

I don't see what they did, I only see what they looked at. And that's not knowledge yet. Just because [Expert 1] clicked on something doesn't mean he understands it now.

Expert 1:

I don't see any application for that myself. But – that was the comment from [Expert 2]: If you know what the person manipulated and what they did, that might be more interesting. We also have this logbook in Mont Terri where we record what work we are doing on an experiment. We also have the option of selecting the object and seeing who last manipulated it. We don't do that consistently. We assign experiments 100 percent, but we don't do other things as consistently. If you were to do that now... But that mainly affects the companies, who say, "Now I've done this and that on this cabinet." Unfortunately, we don't do that consistently or don't have the time to implement it. But the information would be interesting. So if I'm now at the experiment – you showed that nice cube – and I click on it, I can see what has been done in this experiment. That's the database for our logbook. Then, of course, I can filter by company or something like that. I find that even more interesting because then I can find out: Did they do everything? Is anything missing? What was done last? So, I see the connection more in the combination of attaching this logbook to these experiment cubes, for example. But I don't think it's absolutely necessary to see what one person has looked at.

PM:

Let's move on to the last section [of my guide]. My dissertation project is about knowledge management in nuclear technology facilities. Of course, you are not a nuclear technology facility. Nevertheless, I chose you as a test site. Now the question arises as to transferability and applicability. In this context, there are a few interesting characteristics of nuclear technology facilities. The main thing is that you don't want to send people into a controlled area with radiation for longer than necessary. You want to keep the time spent on site as short as possible. Perhaps this can be achieved with an XR application like mine? That would be the case if I could do certain things faster and more efficiently than without glasses.

Expert 2:

So, based on my limited experience with these glasses... When I compare this with my knowledge of the laboratory, I am definitely slower with the glasses. No question about it. Someone who is unfamiliar with the lab but knows the glasses well might be faster. But if I imagine a nuclear facility with a radiation protection area, I would do the whole thing in VR, not in XR/AR. Because then I don't have to go there at all, and I can only work on site with the glasses to a limited extent. If you take a technician who knows their equipment and has to manipulate something, they don't need glasses because they know what they're doing. And if they don't know their way around, then they should use VR, practice at home, and then go in there after they've done it three times virtually, and then you'll probably be faster than if you did it with AR. But that's just my assessment.

Expert 1:

I agree 100%. In most cases, we are allowed to be anywhere. There are a few cases where you could say that the time could be reduced. But if it's complicated, I would rather use VR to learn about the location, the lab, and I would also need longer with the glasses than without them.

Expert 2:

But for me, the use case is not primarily for people who have to do something dangerous or complicated. For me, the use case for us is at least more about obtaining additional information while viewing, walking through, and learning about something. And that's why I find your use case at a nuclear facility... If you send people in there who just want to look around, that's more of a case that doesn't happen. You don't just send spectators into a radiation protection area, that's completely nonsensical.

PM:

OK, I take it that VR would be better than XR/AR here.

Expert 1:

Well, I see it more as a case of displaying essential information that is relevant to safety. I assume that the information is also available on the objects in the nuclear technology facilities. But maybe at some point you become immune to these signs. With the application, I could use warnings to point out the dangers. But it could also be irritating. It takes the focus away from your current work. You've been trained, you know what you have to do, and you want as few distractions as possible. That's why I'm torn on this. If the information could also be displayed when you're not on site, that would be very good.

Expert 2:

[Expert 1] just gave me this idea, where I think it would be useful if you were in a mixed zone where you let people who work there and visitors in. You could set very clear boundaries: "You can't go any further here." For me, that would be almost the only area of application I can see. Yes, if you already have people in there and they're wearing glasses like that, then you can restrict them, so to speak. But these restrictions also exist structurally and visually on the floor and everywhere else. That would be double the effort.

PM:

Let's imagine a nuclear power plant that is being decommissioned. Then you have a different system status every day. Elements are removed, scaffolding is temporarily erected, and usually various trades are working on site. This means that you have very few of your own staff and a lot of external personnel. My concept is a kind of platform that is suitable for integrating external companies. It remains to be seen whether this is done more with VR for planning purposes, i.e., discussing specific work processes in detail, or whether the application is actually used on site. So perhaps my concept is less exciting for the operator on site than for the exchange with external personnel coming from outside.

Expert 2:

If you take that idea further, I could imagine that you have a team of people working there, dismantling, and you have your mandatory safety officer sitting there, standing there, wearing these glasses, and then the daily status is pressed onto his or her eye. That way, the person could guide their employees, as they would see real-time information about the status of the construction site in augmented reality. But that would be a person who doesn't work and basically directs the others.

PM:

That could be the SiGeKo, the safety and health protection coordinator.

Expert 2:

If they wore something like that on a demolition site and could protect their employees from going into the wrong areas or touching the wrong materials... Yes, of course, that's conceivable.

Expert 1:

I don't know enough about the requirements, but anyone I don't need to have on a construction site like that is a good thing. Maybe you could use multiple cameras to record everything and then immediately warn someone if they leave the area where they are working. You could probably do a lot more with AI and so on . But I see the case in training, i.e., getting to know everything in advance. You say the situation is changing. That means the place will soon look different again. And I think it makes sense to be able to see what it looks like from the outside without having to go there. You also had this marking function, for example. So the company can mark things in advance that need to be done. They can look at it together, clarify the work processes in the virtual space before they start discussing on site. I can well imagine something like that. But I can also imagine – if you say that there are scaffolds – that people wearing such glasses are not allowed to walk around on scaffolds and stairs. Because perhaps the current technology does not yet allow you to see everything. So it may well be that it is not permitted. Or if you say, "This is necessary now," but then you have a technology that can sometimes produce errors, that can cause an offset, and so on. So I can also imagine that it might not be permissible at all. But in onboarding, in getting to know each other, it's obviously a brilliant technology. So overall, I don't know enough about the requirements, but I can imagine that there are areas of friction.

PM:

Yes, I think there definitely will be. Next week, I have an expert interview with someone from a nuclear power plant. These are exactly the questions I will be asking.

We have now come to the end of our interview. If you had to sum it up in a nutshell, what do you think are the biggest hurdles and the biggest benefits of my overall approach?

Expert 1:

The biggest benefits, if it is programmed in such a way that it is also optimized in terms of time, so to speak, are ultimately the time savings and the fact that I can see the information on site, even for objects that are not visible. I see really great potential there. The biggest hurdles are, of course, technical again. We talked about georeferencing, that it's not 100% accurate at the moment and that it might take more time because not everything has been programmed to the extent that the application saves me time. And, of course, I can see that some of our colleagues might have difficulties using something like this. On the one hand, perhaps technologically, but on the other hand, also because they know this lab so well that they tell themselves it won't save them any time. And will our external companies ever use such technologies? I see that in the distant future [laughs]. I see great potential for visitors if you can display information in a playful way that would otherwise not be available and show the 3D complexity with a well-designed application. It is, of course, an enormous benefit for a laboratory like this if visitors can immerse themselves in it and better understand its complexity. The spatial aspect in particular is very complicated. The sequences of events in the experiment are difficult to convey. I see great potential here.

Expert 2:

I feel the same way; I would agree with everything that has been said so far. I think it is not only very helpful for visitors, but also for the experiment leaders and for us, i.e., for borehole planning and experiment planning. Sometimes even to understand a borehole, to understand diffusion paths in 3D. 3D models are one thing, but when you stand there and see what it looks like in the ground and how close everything is to each other, I find the extremely helpful. Also, the fact that you stand there together, discuss and talk about the same thing. You can look at things you wouldn't otherwise see, look at them together, as we did in this multiplayer mode. We showed and explained things to each other, and I find that helpful.

From a site management perspective, it's helpful to get an overview of part of the infrastructure, be it electrical, network, or an experiment, things like that. Unfortunately, all the hierarchies, who is connected to whom, are not yet implemented on the BIM side. But if you can display such things naturally, so that I can see at a glance who is affected by a power failure, for example, that would be great.

For me, the downsides are mainly that you can no longer wear a helmet. That's obviously terrible for my safety concept if I have a bunch of people without helmets [laughs...]. The second thing – perhaps only half-serious – is the strain on the eyes. It's not sharp enough for me, the resolution isn't good enough. I find it visually tiring after a relatively short time. But I'm confident that it's only a matter of time before that improves. It's not in the nature of things, it's simply down to the hardware. Nevertheless, I was really impressed with how good it already is. I also think you've done an extremely cool job. You can argue about the whole UI—there is room for improvement—but hey, I see worse every day on my computer.

I also don't think we'll be able to get external staff on board quickly enough for it to add value for them. What would be really great is if you could call up fixed scenarios such as a lab check or a picket assignment as a scenario. Then I could put on the glasses, say "lab check," and off we go. That would be extremely cool. Or you could just say "Visit for geologists" and then it would take me to ten locations and I could look at them. It wouldn't even have to guide me, but then only those things would be there. I would find that cool.

If you have multiplayer scenarios like that, you have to make sure you can pair up faster. The two of us have now restarted the application five times. It didn't bother me, but if I'm visiting with ten people, it'll take me 20 minutes until we've synchronized eight of those ten.

PM:

Sure, it should work the first time.

Expert 2:

Yes. The whole master-slave mechanism and logging out and back into the application with so many people would be a nightmare scenario. It doesn't work that way, but I mean it was a proof of concept for two people and that was good. But I think if you wanted to do that with groups of visitors or something, then this step needs to be easier.

PM:

All right, thank you very much.

Expert 1:

Thank you very much for the interview. You have done a really great job for us. These are super interesting insights. We can gain a lot from this, which is also really relevant for us. And especially the fact that all this BIM data we have can somehow be made even more accessible in 3D space. I think it's super clever.

PM:

Thank you both, see you soon.