Designing for Visualization in Motion: Embedding Visualizations in Swimming Videos

Lijie Yao, Romain Vuillemot, Anastasia Bezerianos, and Petra Isenberg

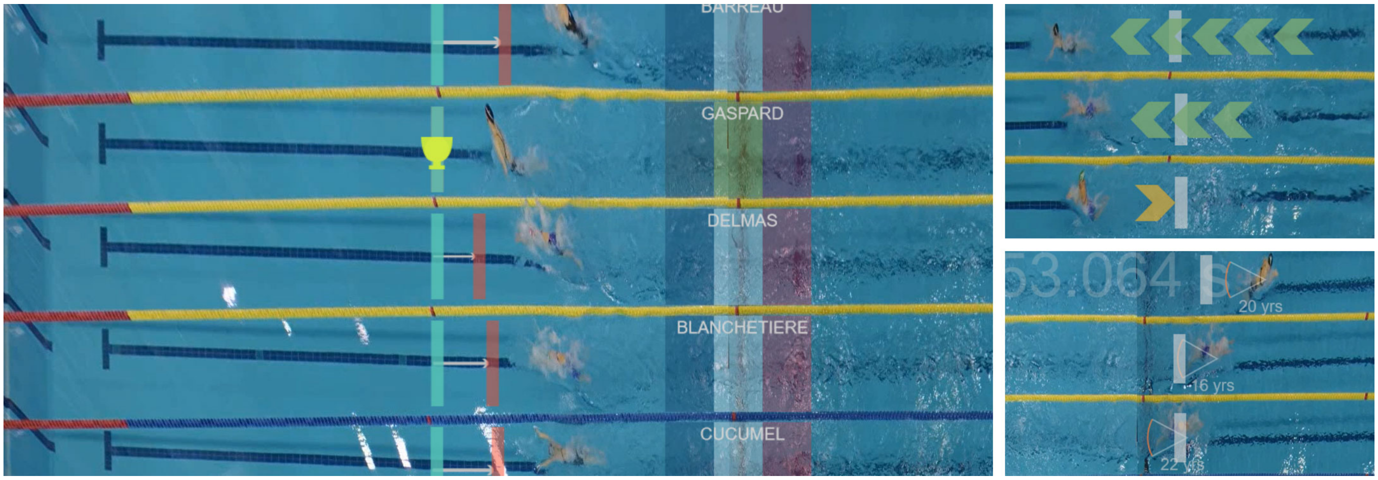

Fig. 1: Embedded representations added to a swimming video of the 2021 French Championship using our technology probe. These show dynamically updating visualizations that move with the swimmers: distance to the leader and predicted winner (left), speed distance to a personal record (top right), and current speed and swimmers' ages (bottom right). The left and bottom right images also show stationary embedded representations of the swimmers' names, nationality, and elapsed time.

***Abstract*—We report on challenges and considerations for supporting design processes for visualizations in motion embedded in sports videos. We derive our insights from analyzing swimming race visualizations and motion-related data, building a technology probe, as well as a study with designers. Understanding how to design situated visualizations in motion is important for a variety of contexts. Competitive sports coverage, in particular, increasingly includes information on athlete or team statistics and records. Although moving visual representations attached to athletes or other targets are starting to appear, systematic investigations on how to best support their design process in the context of sports videos are still missing. Our work makes several contributions in identifying opportunities for visualizations to be added to swimming competition coverage but, most importantly, in identifying requirements and challenges for designing situated visualizations in motion. Our investigations include the analysis of a survey with swimming enthusiasts on their motion-related information needs, an ideation workshop to collect designs and elicit design challenges, the design of a technology probe that allows to create embedded visualizations in motion based on real data, and an evaluation with visualization designers that aimed to understand the benefits of designing directly on videos.**

***Index Terms*—Embedded visualization, sports analytics, design framework, visualization in motion.**

I. Introduction

L. Yao and P. Isenberg are with Université Paris-Saclay, CNRS, Inria. E-mail: yaolijie0219@gmail.com, petra.isenberg@inria.fr.

A. Bezerianos is with Université Paris-Saclay, CNRS, Inria, LISN, 91190, Gif-sur-Yvette, France. E-mail: anastasia.bezerianos@lri.fr.

R. Vuillemot is with Université de Lyon, École Centrale de Lyon, CNRS, UMR5205, LIRIS, F-69134, France. E-mail: romain.vuillemot@ec-lyon.fr.

Manuscript received 27 July 2023; revised 5 December 2023; accepted 9 December 2023.

MANY sporting events embed visual representations of statistics related to athletes' performance and race/game metadata in their broadcasts. These embeddings can range from simple metadata such as the current game/race time to complex computed statistics of shot performances, probabilities of hitting or winning, or speed of balls or other sports equipment. Visual representations often take the form of text, inserted temporarily at fixed locations on the display. With the availability of computer vision methods to track athletes, however, it becomes possible also to create situated *visualizations in motion* [1] that move entirely with the athletes [2] or equipment (e.g., balls, cars, ...). In swimming races, for example, record lines move along the pool according to an invisible athlete swimming at average record speed [3]–[5]; or in basketball [6], [7] visualizations on player shot-probabilities are shown with bar or donut charts above players' heads. Yet, creating and testing designs for visualizations in motion is difficult as real-world sports contexts contain busy backgrounds with various motion characteristics and need to be designed to be informative but not distracting from the audience's primary motivation: seeing and following the athletes and their performance.

In order to explore how to support design processes for such visualizations in motion, we focus on one specific sport, swimming, for the following reasons: (a) Similar to other sports, swimming has rich, dynamic data, that is already visualized to a limited extent. This indicates that the general public may be open to the addition of sophisticated visual representations. (b) There is the first evidence that people can very accurately read quantities from simple moving visualizations [1]. (c)

Working with sports videos is always difficult due to copyright constraints that may make sharing research results difficult. We have a standing collaboration with national swimming associations in our country and, therefore, the right to record, track, and share our own data from national competitions.

Our work involved three phases:

Data Needs Identification & Visualization Ideation: To seek opportunities for embedding visualizations in motion in live swimming races, we first reflected on previously visualized data items in professional competition broadcasts. We found that current visualizations in swimming competitions are very limited to simple data and representations. To further understand which other metadata the viewers of swimming broadcasts would be interested in, we conducted an online survey with 80 swimming enthusiasts. We found that our respondents were highly interested in several data items that are not currently part of professional swimming broadcasts and that would require detailed tracking and metadata collection. Next, to elicit possible designs for these data items and study ideation in this context with traditional methods, we conducted an ideation workshop that involved participants sketching on transparent films to simulate motion effects [8]–[10]. We collected various sketches and noted difficulties in designing for the motion context with the simulated motion. In particular, a lack of complexity and realism of the actual moving video context and non-updating data was mentioned.

Technology Probe Development: To bridge the gap between designing with a motion context and the lack of realistic motions and data updates, we implemented a technology probe [11] that lets users embed, design, and modify visualizations while having access to instantaneous previews. The technology probe includes a playable video with real tracking data, swimming data and visualization selectors, design parameter controls, and a layer panel. Our technology probe is designed to be field tested to inspire users to reflect on designing for motion and to elicit features for future tools. To avoid copyright conflicts, we collected our own videos and tracking data using our own annotation tool.

Evaluation: To investigate the benefits of real-time visual feedback during the design process of embedded situated visualization in motion and elicit feedback on features for future technologies in this context, we conducted a user study with designers. All our studies were preregistered on OSF; supplementary materials are available on https://osf.io/nxyr4/. The technology probe *SwimFlow* can be accessed at https://motion.isenberg.cc/Swimming/index.

Following this research process, the main contributions of our work are findings on how to support the design process for moving situated embedded visualizations in videos. Specifically, we give reasons for why a full-motion design context is important. A secondary contribution involves insights into appropriate data and visualizations for swimming videos.

II. Related work

Our research is closely related to the topic of situated and embedded visualizations [12], which are visualizations that represent information close to a data referent [13]. Specifically, we focus on embedded visualizations in a context involving motion. Our own past work introduced visualization in motion as a concept, a research agenda, and first evaluations of motion features that may affect the readability and design of visualizations in motion [1]. Nevertheless, this prior work did not consider the design *process* and challenges in the design of such visualizations and did not investigate a concrete design scenario as we do here. Here, we work concretely in the area of sports visual analytics and swimming in particular as a use-case. Consequently, we review how visualizations have been explored, designed, and embedded in videos. We complete the related work by discussing relevant related authoring methods.

A. Sports Visual Analytics

Situated visualization has gained a lot of traction in the augmented reality community. In their SportsXR work [14], Lin et al. provide several case studies of situated, sports-specific visualization designs in immersive environments for training, coaching, and fan experiences. Like us, the authors argue that videos play a central role as a situating context for visualization but that the data remains a challenge to extract accurately. Videos also play an important role as a reference [15] and as a validation mechanism since displayed data can be compared with the original scene. VisCommentator is related to our work in that it explored the combination of visualizations and computer vision with sports videos. The tool can, based on user selection, embed both static and animated visualizations in videos [16]. In contrast to our work, the visualizations are animated but not in motion according to our definition. Early work on soccer videos [17], maps players' controlled zones to the soccer field, but only in single video frames.

Several companies (like Footovision [2] or SportsDynamics [18]) commercialize video augmentation tools that allow to embed simple statistical graphics in sports videos, often for highly funded sports such as soccer or basketball. In contrast to our work, the underlying technology is not grounded in empirical work and often focuses on simple effects and highlighting. Visualizations in these tools seem inspired by other video-related domains, such as video games [19]. Here, tiny visualizations, like health bars, are often attached to game characters and move with them. The research work most closely related to ours is embedded basketball visualizations for in-game analysis [7], [20]. In contrast to our work, Lin et al. [7] focused on studying how viewers would control the visibility of the visualizations and iBall [20] helped casual fans understand the game. Instead, we focus on how to support the design process of the visualizations and their embeddings.

Our work is also related to efforts in the visualization community that provide a novel lens on various sports data. Perin et al. [21] summarized designs that demonstrated the potential of sports visualization for both narrative and analytical processes. Most sports have been visualized, even if only little data was available. Examples include tools for table tennis [16], [22]–[24], soccer [17], [25], [26], basketball [27], [28], badminton [29], [30], and tennis [31]. Unlike us, this past work mainly targeted professional coaches and athletes.

We complement this past work by studying how to support the design process of creating embedded visualizations in motion based on real data needs and videos.

B. Authoring Tools for Adding Visualizations to Videos

Authoring tools for visualization allow to create bespoke visualizations in a graphical user interface without programming. In their critical reflection on authoring tools [32] Satyanarayan et al. review three tools: Lyra [33], Data Illustrator [34], and Charticulator [35]. These and other tools focusing on powerful and creative data to mark mappings and configurations [36], [37], are built around the creation of static, non-moving visualizations. While we also provide authoring components in our technology probe, we focus on the discovery and implementation of features needed to design for the situated and moving context in swimming videos. Our probe is thus not as flexible and powerful as the fully-fledged past approaches for authoring bespoke visualizations.

Besides work on GUI-based authoring, researchers have also explored other novel methods. Sporthesia [38] is an example authoring system that takes a language-driven approach in the sports context where visualizations are automatically created based on commentary on the sports activity. Work in Augmented Reality (AR) has looked at how to design visualizations for a dynamic context. For example, RealitySketch [39] can attach simple statistical charts on a video based on tracked objects. The sketched charts relate to spatial data extracted from the object tracking. MARVisT [40] is an authoring tool for the general public. It embeds glyph-based visual representations in AR by binding the glyphs to objects in the environment.

More broadly, our work relates to work on data videos. Data videos attempt to tell a story with data visualizations. Much of the literature centers around questions of how to create visualizations in videos to form a coherent narrative and an enjoyable watching experience [41]. Some work, however, has also considered how to embed charts directly in videos [15], [42]. In their design guidelines, Tang et al. [15], [42] mention motion factors, but these are primarily related to animating visualizations. Yet, several of their design and data-related considerations are relevant to us even though the authors considered non-sports-related videos. The authors recommend using colors that are in harmony with the videos, appropriate visual mappings, placing visualizations next to objects in the video, and avoiding overlap. Their design goals similarly apply to our context: avoid conflicts between visualization and video intent, enhance perception, increase appeal, reduce cognitive load, emphasize the data, and keep consistency. Ultimately, our goal is to improve our understanding of design processes that best allow designers to juggle these important factors in the creation of embedded visualization in motion.

III. Data Exploration & Visualization Ideation

In order to base our exploration of design processes for situated visualizations in motion in current practices and real data needs, we followed a three-step process. We first analyzed current data and visualizations embedded in example swimming races. Our findings were then inspired to inform the creation of a crowdsourced survey that we ran to explore what an extended set of data swimming audiences might be interested in. Our survey was pre-registered at https://osf.io/qdbhg. Finally, we conducted an ideation workshop to generate designs corresponding to the data items that were most interesting according to our survey. Our goal was to elicit designs but also to see how participants would fare with a traditional design method in this context.

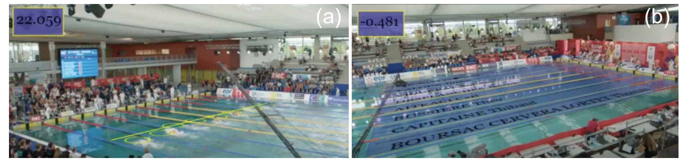

Fig. 2: Recreated examples of embedded visualizations for swimming from an internal prototype. (a) Swimmers' current position circles and record line moving with the swimmers' movement. (b) Stationary swimmers' names in each lane. (a)(b) have a stationary timer in the top left corner. (Due to image permission considerations, these examples are from our own corpus).

A. Review of Current Practice

To understand how to best support the design process of embedded visualizations, we first studied how visualizations are currently embedded in swimming broadcasts from three aspects: (a) which data is displayed using which representation, (b) visualization movement and placement on the screen, and (c) under which camera positions and perspectives. To explore these embedded visualization practices, we looked at a large corpus of videos, including the Olympics from 2008—2020 as well as the latest FINA world championships and the French nationals in their latest instance. We chose to focus on the 2020 Olympics races because they used more embedded representations than any of the other competitions.

The videos we analyzed covered all 4 strokes (freestyle, backstroke, breaststroke, and butterfly) and all 5 race lengths (50m, 100m, 200m, 400m, and the 4 × 100m medley). We looked at each race from the start signal until the last swimmer arrived. The detailed analysis is in the supplementary material.

Visualized Data and Representations: We classified the data visualized in the current races into swimmers' metadata, including nationalities, names, and lane numbers; temporal information, such as the time taken from the start of the race to the current time (elapsed time) and the lap time difference between swimmers and a specific record; record-related information, like the word and the competition record; speed-related data — current speeds; and distance-related data — distance swam. Apart from the nationality represented by a flag and the record shown by a colored line (Fig. 2 (a)), all other data was displayed in text (*e.g.*, Fig. 2 (b)).

Reflection: These data can be grouped into `dynamic` data, whose value changes over the course of a race (*e.g.* speed), and `static` data, whose value stays consistent (*e.g.* nationality). We noticed that the current visualized data covers only a very small part of the data that can be collected about swimming races. For example, the distance between two swimmers can be easily calculated according to the distance swam, and it is possible to predict if a

TABLE I: Swimming data matrix with example data items. A yellow highlight indicates the data items that are visualized statically in current swimming races; purple highlights the data captured and shown in motion. Elapsed time is the time gone by since the beginning of the race. Lap time is the time taken to complete one length of the pool. History of speed is a continuous time series representing the speed changes in the race. A passing happens when a swimmer overtakes another swimmer. A stroke is an arm movement with a specific technique used to propel the swimmer through the water.

Dynamic Data					
Time related	Elapsed time	Current lap	Average lap	Lap differences to other swimmers	Lap differences to a record
Speed related	Current speed	Average speed, Speed history	Acceleration, slow down	Speed differences to other swimmers	Speed differences to a record
Distance related	Distance swam	Remaining distance	Distance differences to leader	Side-by-side distance between two swimmers	Trace of movement
Predictions	Record break	Winner	Completion time	Next passing	
Swimming techniques	Distance per stroke	Reaction time	Diving distance	Stroke count	
Static Data					
Swimmers' metadata	Nationality	Name	Lane number	Age, Gender	Height, weight
Record related	World record	Competition record	National record	Personal record	
External	Social Media followers	Social Media discussion	Sponsor		

swimmer may break a record. To more broadly cover potential data to be visualized, we first created a data matrix with the categories identified in the videos and reported before (5 data categories and 10 data items). We then expanded the data matrix through discussions with swimming practitioners and swimming-focused researchers who know of other data that can be captured during races and that they expressed particular interest in (an additional 3 data categories and 27 data items). The final data matrix can be seen in Table I. It includes 8 categories (with 37 example data items): time-related, speed-related, distance-related, record-related, external data, predictions, swimming techniques, and swimmer's metadata.

While we identified a large number of data items that could be visualized, which data might interest the general audience remained unknown, and how best to visualize them in the context of swimming races was similarly unclear. Therefore, we followed up with an online survey (Section III-B).

Movement Status and Situatedness: The majority of the currently visualized data in swimming videos is stationary without a change position on the display. Exceptions were speed labels and record lines: When swimmers were close to the end of their lane, a speed label, composed of the text of their current speed and a flag for their nationality. Record lines (Fig. 2 (a)) moved at the record holder's average speed.[1] These two moving visualizations were embedded close to the swimmer or embedded in the swimming pool (Fig. 2 (b)). The remaining visualized data remained static on the screen, usually in corner positions, with a large distance to the swimmer they were referring to. (*e.g.*, Fig. 2: timer).

[1]A record is recorded as the time taken to complete a certain distance.

Reflection: We saw that compared to the early Olympics videos (2008) we had initially looked at, the 2020 videos we analyzed included a more diverse set of embedded visualizations. We also saw that existing visualizations became more detailed, for example, by adding labels to explain the visualizations themselves. This might indicate an increased interest of audiences and public broadcasters to see and show embedded visualizations in sports. Yet, how to design future visualizations, how to attach them to a moving reference (*e.g.*, a swimmer), and how to define their movement paths remains to be explored. In addition, there are currently no tools to let designers easily experiment with different designs. Designers may want to experiment with different visualization placements around a moving target, testing that the visualization remains legible throughout the moving trajectory, and checking if the data update frequency is too slow or too fast and, thus, distracting. In order to explore features of future prototyping tools for embedded visualizations in motion, we developed a technology probe [11] we call *SwimFlow* (Section IV).

Camera Shots: How a camera points at the swimming pool is important for the rendering of embedded visualizations as different camera shots influence the size and position of each swimming lane and, with it, the visualizations projected onto them. We found multiple camera shots used in swimming videos. These shots came from cameras positioned around the pool or underwater. The underwater shots were usually only shown briefly to highlight a small number of swimmers, often even only focusing on a single swimmer. Future visualizations shown in underwater scenes could benefit from the pure and consistent blue background of the pool, with very little visual interference (such as a few lane lines). The shots from outside the pool were more common. They showed overviews of multiple, if not all, swimmers and, as such, might be better for including comparative visualizations. On the other hand, these shots come at the cost of a more noisy and colorful background context to place the visualizations in (multiple lane lines, audiences, reflections, lighting, *etc.*). The camera shots outside and around the pool included a bird's-eye view (top), side views, diagonal views, and transitions between them. In the diagonal view, swimmers swam along the diagonal of the screen – from the bottom right to the top left or from the bottom left to the top right corner.

Reflection: Similar to many other sports, swimming broadcasts change shots frequently and also include moving cameras that pan and rotate. For example, when swimmers turned around it was common to see a switch from an in-air shot to an underwater shot focused on the current leader. Consequently, not only did the swimmers' positions change on the screen, but also the swimmers' sizes and the angles from which they were shown could change within a few seconds or even less. For embedding data in swimming broadcasts, this means that viewers will have to be able to track visualizations across shots and be able to deal with the complexities of changing backgrounds and motion added by camera movement.

B. Surveying Data Interests

As we saw previously, despite the many types of data tracking that may be of interest to audiences, currently

TABLE II: Participants' swimming race familiarity: distribution of races watched and their watching frequency.

	Olympics	International	Continental	National	Regional
All races	15	8	5	3	1
Almost all races	29	15	13	12	4
Some races	32	33	33	29	15
No races	4	24	29	36	60

visualized data in swimming races are quite limited, and embedded visualizations in motion are even more rare (Table I). There might be many reasons behind this scarcity of embedded visualizations: designers might not be aware of what audiences want to know, they might be afraid of causing too much distraction, or they might not have the data available. The first two challenges can be addressed by visualization research, the latter is more of a problem for computer vision and dedicated prototyping tools. Here, we start by providing a first empirical investigation of the data needs of swimming audiences and then move on to discuss how to support the design process.

Survey Procedure: We conducted an online survey on LimeSurvey [43] to gauge general audiences' interest in specific data while watching swimming competitions on TV or live stream. We advertised the survey to swimming associations, on social media, and at a poster presentation [44]. We collected basic demographic information and the frequency at which participants watched certain swimming races on TV. Since the visual encodings embedded in broadcasts might differ by region, we asked participants to report their broad geographic location. The main part of the survey asked participants to rate their interest level per data item (Table I) on a 5-point Likert scale from not interested at all to extremely interested, with an extra option "I do not know/I did not understand the question." To avoid people choosing the last option, we gave participants explanations for each data item in text or graphical form (e.g., graphics for distance difference to leader & side-by-side distance between two swimmers). We also allowed participants to add additional data items of interest. The order of data categories, as well as data items, was randomized per participant. During the description of results, we make use of the data type icons from Table I and the representation icons in Section III-A. For example, a static representation of metadata using a symbol would be described as : × [].

Participants: In total, we gathered complete answers from 80 participants: 27/80 ♀, 52/80 ♂, 1/80 unspecified; 65/80 lived in Europe, 8/80 lived in North America, 6/80 lived in Asia, 1/80 lived in Africa; the ages of participants ranged from 18 to older than 75 [2]. Participants' familiarity with swimming races is reported in Table II. 75/80 participants reported that they had already seen visualizations in motion on TV in the form of nationality flags (: × []), current speed text (: × []), and record lines (: × []).

Results and Findings: Participants' interest in each data item is depicted in Fig. 3. In blue on the right is the percentage of participants who expressed interest. Labels for data items with an interest level over 70% (30/37) are shown on the left, and below 70% (7/37) on the right. The data items added by participants are not shown in the figure because they were rare but can be found in the supplementary material.

Data Categories of Interest: Out of the data items that received an interest level above 70%, 20/30 belonged to dynamic updating data, while the remaining 10/30 were static data. Participants found all time-related, speed-related, predictions, swimming techniques, and record-related interesting, while external data received the least amount of interest.

Data Items of Interest: The three data items we found moving with swimmers in the past Olympic broadcasts (current speed, flags, and record lines) received high-interest rankings (above 89%). The world record in particular was ranked as extremely interesting by participants. However, we also saw data items at the top of our list that were not yet part of current broadcasts. The distance between the current leader and other swimmers stood out with the second highest rating.

The participants also expressed interest in other subtle differences between swimmers, including time-related and speed-related data, such as lap time differences and speed differences to a record and/or other swimmers. Of these, we only saw lap differences from the world record be briefly displayed in races at the national level or higher.

In addition to the differences between swimmers, the participants showed a keen interest in the swimmers themselves. This makes sense since being submerged in water makes it difficult to identify who is who. Participants highly valued swimmer's metadata, such as their names and personal records. Surprisingly, our respondents were also highly interested in swimming techniques, including the swimmers' reaction time and the diving distance. This data is not only absent from the current Olympic broadcast but also rarely orally mentioned by commentators.

Some of the data that was the least interesting we had considered interesting ourselves: swimmers' movement traces, for example, and the distance to neighboring lane. Gender was also not of much interest, perhaps because the mixed-gender medley only made its debut at the 2020 Tokyo Olympics.

In summary, our survey showed a huge potential for augmenting swimming videos with embedded data about the race. Only 7 of our 37 data items had interest levels below 70%, and we found many data items of interest that are not yet common in public swimming broadcasts.

C. An Ideation Workshop

The main goal of our following ideation workshop was to elicit possible representations for popular data items for our technology probe. In addition, we observed challenges of traditional sketching as an ideation technique [45] for a motion context. The participants of our workshop were 4 senior visualization researchers (including two co-authors) and 3 students (2 Masters and 1 PhD student) working in the area of visualization. The workshop was led by the first author of this paper, who did not participate in the exercise.

[2]distribution of age range: 18-24, 25-34, 35-44, 45-54, 55-64, 65-74, >=75

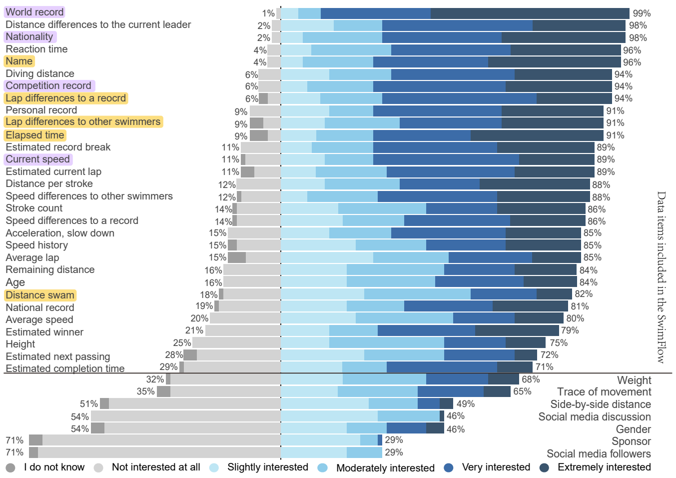

Fig. 3: Survey results: Participants' interest level in seeing visualization in swimming races for each data item. A black horizontal line separates the data items included in our technology probe (Section IV). A yellow highlight indicates the data items that are already visualized statically in current races; purple highlights the data visualized in motion.

Fig. 4: Workshop procedure: A participant (a) sketches her design on a transparent sheet with a printed background below, and (b) presents her design by attaching the transparent sheet on a whiteboard for all others to see.

Procedure: The facilitator first presented a summary of the review (Section III-A), as well as the results of the online survey (Section III-B), and introduced the sketching materials. Next, participants completed one design round per data category (Table I). In each round, participants had 10 minutes to sketch a design (Fig. 4 (a)) and 3 minutes to present (Fig. 4 (b)). Participants sketched their designs (Table I) on transparent sheets placed over one of five printed random frames from the women's 200m butterfly final at the Tokyo 2020 Olympics [46]. Participants could slide their transparent sheets over the background images to simulate motion. At the end of the workshop, participants voted for their favorite designs.

Designs and Results: We saw a variety of designs (n=46, all available in supp. materials). People drew representations already seen in broadcasts (nationality flags, lines, and text) but also a variety of new representations. This is not surprising as all participants had a visualization background. We grouped the sketched visual representations into 4 categories: `graphics` (21/46) were small pictogrammatic, iconic, or symbolic graphics that may indicate positions and categorical information. They could also be custom data-driven graphics akin to data glyphs that change shape based on data. Examples included flags to show who is predicted to win or arrows whose length and width encode the distance to a leader; `lines` (14/46) were representations of absolute or relative positions drawn to represent, for example, record lines, position traces, or two parallel lines whose position indicates a distance difference; `text` (15/46) was used to write out quantitative or categorical information; for example, the current race time or the name of a swimmer; and `charts` (11/46) were traditional data representations such as pie, donut, or bar charts. 32/46 visualizations were designed to connect directly to the `swimmers` and move with them, while 14/46 visualizations were located relative to the `swimming pool`. Participants tended to embed simple data items, such as the current speed, as close as possible to the swimmer, sometimes even overlapping with the swimmer. Participants preferred to attach more complex data items, which require more display space, to the swimming pool even if the data was dynamically updated. For example, one participant attached a visualization about the estimated winner and three participants embedded distance swam/rest to the side of the pool. 16/46 visualizations combined multiple representations, such as dynamically updating data with annotations and comparisons between/across swimmers.

Sketching Challenges: We had prepared the sketching exercise in a way that would allow people to experiment with motion using transparent overlays akin to sketching techniques for user interactions [47]. In practice, we observed that this did not work very well for the participants. Motion factors could only be tested well when the underlying images of the pool included no perspective distortion; when a design was overlayed on a pool screenshot taken at an angle, the design would slide into different lanes when the transparent sheet was moved, and the fact that the design did not correctly change size became disturbing. Moreover, moving components that could affect visuals, such as the visible part of the swimmer, splash, and shadows, could not be accounted for. Besides, the movement effects of sliding the transparent sheets are only a part of the factors that need to be taken into account — the asynchronous motion between background, data referent, and visualizations made it additionally difficult for participants to imagine what their designs might look like, especially when the designs showed dynamically updating data.

IV. Towards In-Context Editable Vis in Motion: A Technology Probe — *SwimFlow*

To be able to study how to better design for the motion context, we implemented a technology probe — *SwimFlow* that included a set of basic features for prototyping visualizations in motion coupled with a video. Technology probes [11] are simple, flexible technologies used to field test the usage of technology in real-world settings. They are not prototypes of fully functional systems but instead are tools to inspire ideas for new future technologies. As such, our probe includes the basic requirements for deployment with designers: real data, an underlying video, a simple set of visualization authoring features, features to define the embedding of visualizations, and video playback options. In this section, we first describe the design of *SwimFlow*, followed by the process of preparing video and tracking data, and end by introducing its components.

A. Authoring Probe Design

SwimFlow (Fig. 7) targets professional or amateur designers who regularly use prototyping tools as part of their workflow. *SwimFlow* allows users to synchronously embed visualizations, edit designs, and preview motion effects corresponding to the moving entities in a playable video. While we identified the need to support multiple camera shots in our exploration of current practices, for our probe, we decided to use only a single shot. Including multiple shots would have required us to have a much more complicated tool and data backend than we needed to study the prototyping of situated visualizations on video. As such, we chose to provide a bird's-eyes view that gave a good overview of the entire pool and all swimmers and a stable visual focus. Thus, any issues found in designing with our probe and camera shot would also affect more complex design settings. *SwimFlow* has three main parts: (a) a demo video with attached tracking data, (b) a set of visualization layers over a playable video, and (c) an interactive UI. The visualizations are rendered as overlays over the video — using tracking data, visualizations can move synchronously with data referents but stay on their own visualization layer. The separation of video and visualization layer allows us: (a) to achieve smooth motion and (b) to make it possible for *SwimFlow* to adapt to various videos and shot types. We developed the interface using web technologies, including JS, HTML, and CSS. The graphical implementation and motion rendering were realized on a Canvas object. We prepared the video we currently use in *SwimFlow* as well as its tracking data ourselves to avoid later copyright issues with sharing our work.

B. Video and Tracking Data Preparation

To have a fixed spatial reference system, the demo video embedded in *SwimFlow* underwent several pre-processing steps, including combining two videos, which were recorded separately from the side stands of a swimming pool (side view: Fig. 6 (a)(b)), into a unique video *as seen from above* (bird's-eye view: Fig. 6 (c)). Here, we describe our video combination, which we realized using a standard computer vision warping technique and our tracking data extraction. All scripts and tools are in the supplementary material.

Video Recording: We recorded multiple swimming races during the *2021 Montpellier French National Championship*, with the authorization of the French National Swimming Federation (FFN [48]). We used the Women 100m breaststroke video in our work.

We set up two side-by-side cameras (Fig. 5), as one camera could not cover the entire pool or would have created too much distortion due to a Fisheye effect. Each camera covered half of the swimming pool and captured the race in 4K resolution at 50 frames per second.

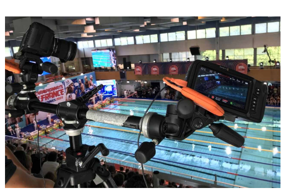

Fig. 5: Video recording setup: left and right cameras independently recorded half of the pool.

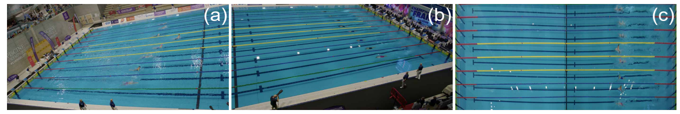

Fig. 6: Screen shots of the self-recorded video from the left (a), the right (b) camera, and the final combined video of a simulated bird's-eye view (c).

Video Combination: To create the bird's-eye view of the entire pool (Fig. 6 (c)), we followed the following process:

- To remove unwanted camera movements, we detected key points on every frame to achieve *video stabilization* (usually the field landmarks, such as the center and the corners on most sports fields or the distance markers alongside the run tracks).
- For two videos to be played simultaneously, we manually selected 4 representative points of each half side of the swimming pool to achieve *homographic projection*. Then, we calculated the video transformations to simulate a camera as seen from above the pool using the OpenCV 4 `warpPerspective` function.
- To reach color balancing and purify the background, we matched color tones and calculated the median image to achieve *background homogenization*.

Tracking Data Annotation: To generate tracking data of swimmers' positions, we manually annotated the video using an in-house annotation tool (similar to labelImg [49]), designed explicitly for annotating swimmers' positions. As annotating swimmers' positions for each video frame is time-consuming, we annotated the *(x, y)* coordinates per swimmer per stroke (*i.e.*, for breaststroke, occurring, in general, less than one second. We annotated positions when the swimmer's head was at its highest). To obtain continuous tracking data for each video frame, we interpolated strokes by multiplying the average speed between two strokes and the delta time per video frame.

C. User Interface Introduction

SwimFlow contains many features. Instead of describing them individually, we group features by function type. Each group of features contains multiple UI elements. Since our goal was to ultimately use *SwimFlow* as a technology probe, we did not attempt to implement features as complex as those present in other visualization authoring systems (see Section II-B).

Setup (S): Designers begin by selecting a data item to represent from a drop-down menu (Fig. 7 (S.1)). The dropdown menu contains all data items with an interest level over 75% from our online survey—together with a donut chart of its exact interest percentage. Next, users can choose from a pre-defined representation we made available based on the designs collected in our workshop (Fig. 7 (S.2)). Other pre-defined visualizations can be easily added to the probe. Once data and a basic representation are selected, the designer can proceed to customize the visualization and its embedding parameters (Fig. 7 (R)) and play it back (Fig. 7 (P)).

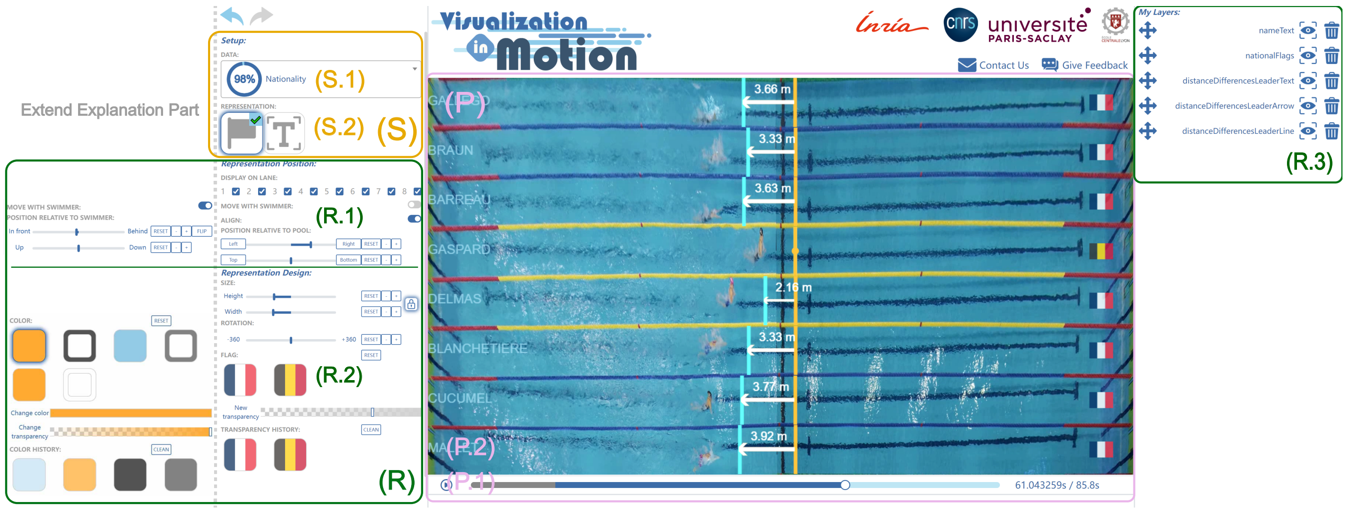

Fig. 7: *SwimFlow*: (S) Setup: Data item (S.1) and representation (S.2) selectors. (R) Embedding parameters - R.1: a lane selector, a switch to attach visualizations to swimmers or the pool, a switch to align visualizations, and sliders to modify the distance between visualizations and their data referents; Visual encoding parameters - R.2: Sliders to adjust size and rotation, a color picker and transparency slider, and layers; Visualization combination - R.3: a layer panel to adjust the visibility and hierarchy of visualizations. (P) Playback: a video play/pause button with a video progress bar (P.1), the video view (P.2).

Fig. 8: Remaining distance indicators once attached to and moving with swimmers (a) and attached to the bottom edge of each lane of the Swimming pool (b).

Visualization and Embedding Customization (R): Core features of *SwimFlow* concern the embedding specification of each visualizations. Designers can select for which lanes the visualizations should be displayed with a *lane selector*; and for which data referents (swimmers or lanes). By default, visualizations are set to be connected with swimmers and move with them. Users can move the visualizations above/below the swimmers or closer/further to the swimmers by dragging position sliders. Users can switch the data referent from swimmers to the swimming pool by turning off the *"move with swimmer"* switch. Once turned off, visualizations will remain at the position where they last stopped. The position controls will now act in reference to the swimming pool (Fig. 8). The *align* switch can then be used to align visualizations and edge buttons can be clicked to embed visualizations at one of the 4 edges of the pool.

SwimFlow also provides the flexibility to flip the moving visualizations' position when swimmers turn. With a pressed *flip button*, a visualization behind the swimmer during their forward movement will automatically flip to be in front of the swimmer after their turn. Additionally, representations with an orientation property, such as an indicated arrow, will be automatically mirrored once the swimmer turns.

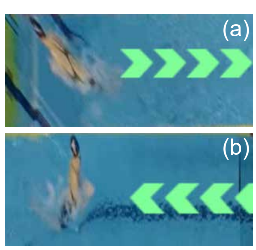

Fig. 9: The position of the acceleration graphic (a) is flipped after their turn (b), and its direction is mirrored.

Representations can be modified in terms of their size, rotation, color, and transparency. Users can manipulate sliders to modify the height and width of graphics, lines, and charts, as well as the font size and stroke weight of text. A color picker and transparency slider allow to modify colors for most elements (except for certain graphics such as flags). When a color is modified, a color history block is created to aid in the reuse of colors.

A *layer* panel similar to those in common graphic editing tools (Fig. 7 (R.3)) allows to combine multiple visualizations. Users can save the current design to layers and adjust the overlapping order by dragging and dropping the layer labels. Users can directly reload an already saved layer by single-clicking the layer name or rename layers by double-clicking. For all lanes users can make a visualization invisible using the visible icon and delete a visualization with the trash icon.

Playback (P): To test how a visualization appears under motion *SwimFlow* incorporates a video play/pause button and a video progress bar (Fig. 7 (P.1)), giving users the ability to not only see the visualizations move but also to get a feel for how the represented data changes over time. The designs can be modified while the video is either playing or paused.

Additional features: *SwimFlow* also includes undo/redo buttons in the top left corner that act on the currently selected data item and actions performed on its current representation. Furthermore, users can share their designs as non-modifiable playable videos via a shareable link offered by *SwimFlow* [50].

D. *SwimFlow Summary*

The SwimFlow design probe has partially similar functions as general video editing tools have (e.g., timelines, layers, undo, redo…). However, SwimFlow additionally has a large number of functions that do not exist in conventional video editing tools. It provides rich design parameter controllers that support users in choosing and customizing visualizations, modifying their size and aesthetics. In addition, SwimFlow allows to embed visualization statically or — based on pre-prepared tracking data — dynamically in relation to moving data referents. And to customize the visualization movement properties.

Although some video editing tools support users in embedding visual designs on video, there is usually no data binding between the visualizations and the data referents. In contrast, SwimFlow supports users in accessing real, dynamically updated data via an internal connection and binding between the visualization, the data, and the context. In SwimFlow, data can be changed in real-time, and visualizations are updated immediately according to the data change. SwimFlow fills a gap in the space of dedicated visualization authoring tools in that it allows to manipulate changing data and instantaneously embed, design, and adjust visualizations in a motion context. Nevertheless, our goal is not to provide a holistic tool. Instead, we developed SwimFlow as a technology probe to help us study the design of visualizations in motion, more specifically to inspire participants to reflect on designing under motion and to elicit features for future complete tools.

V. Effects of Motion Context: A Design Study

To understand the impact of real motion as a context during the visualization design process, we ran a user study with *SwimFlow*. Specifically, we wanted to identify the specific benefits of designing directly inside the video context. As one can perhaps already gauge from the previous sections, future tools for designing embedded visualization in motion will have to be feature-rich and will take considerable effort to develop. Whether the benefits of prototyping in the context of videos will be worth the required development effort is one main question we asked ourselves. As such, we wanted to compare the design of embedded visualization in a static way (e.g., on still images, as is common in rapid prototyping using sketching) to the design of visualizations under motion (e.g., on a playable video, with both data and referents being dynamically updated). Instead of focusing on one specific and a-priori defined video, our goal is (a) to explore more generally the impact of motion context on design and (b) to understand what are the opportunities and challenges of design visualization in motion in a real application scenario. Our participants were 8 graphic designers who had visualization experience and could, thus, reflect on the impact of motion on their process as well as future design requirements.

A. Method & Procedure

We had initially planned and pre-registered (https://osf.io/fw7gj) an online experiment but changed to an in-person one for all but one participant. For the first participant, we conducted a video call but found that the resources needed to run the online session interfered with a smooth design preview experience; we then switched to in-person experiments for all subsequent participants.

The study had two phases. In the first, designers created their visualizations using *SwimFlow* on a static screenshot from the video. After they had completed their design, we showed them what their design looked like on the video. In the second phase, participants refined their design on the video with all features for video playback enabled. The study lasted approximately 1 hour. We collected quantitative and qualitative feedback to understand the differences between the two design processes.

The study began with a consent form and a short demographic questionnaire. We first showed participants an Olympics swimming video to give them an understanding of the context and application scenario they would be designing for. Next, we showed a tool tutorial video and gave a 5-minute tool exploration time in which participants could experiment with the visualization generation and embedding features. Then, participants began Phase 1:

`Motion-limited mode:` We presented a stripped-down version of the tool in which the play/pause button and video progress bar were hidden. Instead, we showed a video frame in which swimmers were at the left-center, center, or right-center of the pool. The frames we used were different for each participant and included different stages of the race (towards the beginning, in the middle, or towards the end of the race). We ensured that in all frames, the space in front and behind all swimmers was large enough to allow room for embedding visualizations. Designers could use all other features of *SwimFlow* related to setup (S) and data visualization (R) but could only preview their design on static frames. In addition, participants could drag sliders that determined the x- or y- position of the visualizations to simulate movement , akin to the sketching scenario in our workshop (Section III-C). There was, similarly, no dynamic update of the visualized data nor any movement of the swimmers. The choice to let our design probe behave as a digital version of the method we used in our ideation workshop was deliberate. As such, it acted as a baseline to study the challenges of designing without dedicated motion support. In the 15-minute duration of this phase, we asked participants to create at least 3 data representations but gave no upper limit. Once the time was up, they were allowed to finish the design of the visualization they were currently working on but not to add new ones. We then asked our participants to rate on a 7-point Likert scale how satisfied they were with their present design. Next, participants visited a link to see a playable video with their designs embedded. After participants had previewed their designs on the playing video, we asked them to rate their satisfaction again with the same rating scale. We then conducted a first interview, asking participants what changes they would like to make to their designs after seeing all the embedded designs in the playing video. Our participants then moved to the second phase of the study:

`Full-motion mode:` In this phase, the video play/pause button and video progress bar were available to participants. They could design on the playing video and simultaneously preview the data updates and all motion effects (moving swimmers, water splashes, . . .). They could also stop the video, change their design, and restart the video again. We asked participants to continue to improve their previous designs and told them that they were free to add/remove/change data, designs, as well as their embedding parameters in any way they wished (with a playing video or without). Once they were satisfied with their final designs or 10 minutes had passed, they could submit their design. Similar to the previous phase, participants were allowed to finish the current ongoing design once they had reached the 10-minute mark. After submitting their designs, participants could again view a playable video of their final result via a sharable link that *SwimFlow* provided. Again, we asked participants to rate their satisfaction with their design on a 7-point Likert scale. We also asked to what extent they thought the final designs had improved and which mode they preferred to design in. We finally asked several open-ended questions on the design process in a second interview, with a focus on asking them to elaborate on their overall preference, what they had envisioned but could not do with the tool, and, more generally, to elicit feedback on missing features.

B. Participants

As we set out to understand the impact of motion in the design process, we focused on participants with existing design experiences that could reflect on their process. We recruited them by advertising via mailing lists, through the network of authors, as well as through contacts mentioned by our participants. In total, we recruited 4 ♀, 3 ♂, and 1 gender-unspecified graphic designer. Participants reported their ages in 10-year bins from 18–54 years; 1 3 3 1 with most participants in the ranges of 25—44 years. They were either professional designers 5/8 or self-reported as "partial designers" (doing design work as a part-time job or having been trained as a designer previously). The non-professional designers were currently either visualization and human-computer interaction researchers 2/8 (including a design instructor) or students in design 1/8. They reported a high level of experience in graphic design (Table III 1st row). Participants also reported using design tools and reading or creating visualizations frequently. Details are found in our supplementary material. We note that visualizations in motion and embedded visualization is a relatively new research direction in visualization [1], [12] and the topic is highly specific. Thus, as of yet, there are no dedicated "embedded or moving visualization designers" in the labor market. Our participants, however, were qualified to design visualizations and have done so in various contexts before. All 8 participants involved in our experiment are professionals or formal graphic designers, 3/8 participants had over 15 years of design experience, while another 2 participants had experience in design with motion factors.

C. Results: Design Decisions Making

We report here the designs produced from our study and the analysis of the first interview.

Designs Produced & Changes Made: While the purpose of this study is not to critique the produced designs, we report a high-level summary as an overview of the features used. In `motion-limited mode`, participants created 44 visualizations, from 4–7 per participant 4 7. The majority of the designs 31/44 were meant to be in motion and attached to `swimmers`, and the remaining 13 were meant to be attached statically to the pool. Visualizations in motion represented data from some of our data categories, including

TABLE III: Study with Designers. 1st row: Participants' self-reported design experience in years. 2nd row: Participants' satisfaction before and after seeing the video in `motion-limited mode` as well as with the final design in `full-motion mode` on a 7-point Likert Scale. 1: not satisfied at all; 7: extremely satisfied. 3rd row: Participants' responses on a 5-point Likert scale about the extent `full-motion mode` to which their design improved. 1: not improved at all; 5: extremely improved. 4th row: Participants' preference between `motion-limited mode`, `full-motion mode`, and `non-preference` options. Red arrows: Participants' satisfaction decreased in `motion-limited mode` AFTER seeing the video; Green arrows: participants' satisfaction increased with `full-motion mode`.

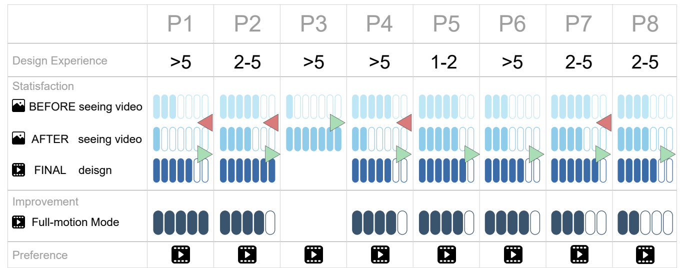

swimmer's metadata, speed-, distance, and record-related data. The most common data items were the swimmers' nationality and distance to the leader, both chosen by 7 times and various representations of world, Olympic, or personal records, as well as the current speed (5× each; × stands for the word "times" here and below). Visualizations in motion were never represented with traditional charts but almost equally often with `text` (12×) and `graphics` (11×). Lines were used 8×. The static visualizations were very diverse. Only the swimmers' name 4/13 occurred more than once. `Text`was the most frequently used representation type 8/13, followed by Graphics (4×) and Lines (1×).

In `full-motion mode`, participants embedded a total of 48 designs, from 4–8 per person 4 8. The majority were in motion 30/48 and attached to `swimmers`. The remaining 18 designs were static and attached to the pool. Compared to the motion-limited design mode, visualizations in motion now came from all our data categories. The distance to the leader was clearly the most common representation 9/30. The number of nationality representations saw the largest drop from 7 to 4. Instead, participants added representations of predicted record breaks (3×). The number of text representations dropped from 12 to 8, while `graphics` became the most frequent representation type (12×) followed by `lines` (10×). For visualizations staying static on the screen, the most frequent representations were of metadata, specifically nationality and swimmer names (4/18 each). The order of frequently used representations for static visualizations was the same as in the motion-limited design mode: text was the most frequent (9×), followed by graphics (7×), and one chart and line representation each.

Designs contained 4 more representations in the full-motion mode than in the motion-limited mode. However, visualizations in motion embedded in the video decreased from 75% to 50% of the total designs because participants found their initial designs too overwhelming and reduced the amount of moving visualizations. 6/8 participants added one or more new representations in the full-motion design mode. Half of the participants removed 1–3 previously added representations. Of the existing designs, 3/8 participants changed their motion status (2 designs were made static, and 1 changed to in-motion). All participants adjusted visual encoding parameters for their representations during the full-motion design phase.

Reasons for Changed Designs: Based on the designers' interview responses, we explain the reasons behind modifications made to designs in full-motion mode.

Added representations: Six participants added representations for the following reasons: (a) to complete information, which no longer seemed to be self-explanatory after seeing the embedded visualizations change on the played video (e.g., adding a `text` annotation for the `lines` of distance differences to the current leader) (2 participants); (b) to compare between data items that had previously not seemed important; for example, one participant added a diving distance representation to compare with the distance differences to the current leader. Their goal was to see if a starting advantage could be maintained over the race (1 participant); (c) two participants found that they visually had more space to embed visualizations in the full-motion mode; and (d) two participants added more data due to a newly found personal interest.

Removed representations: Four participants removed representations. (a) One participant found some data not as useful as previously thought because it had been hard to imagine how data changed over time. For example, one participant found that only one was enough to be shown for a national competition; (b) another participant found that some data would not change over the course of the race, so it was unnecessary to let it move with swimmers as it might distract the audience. The lap time, for example, updates only after one lap is completed; (c) two participants cited changed interests.

Modified representations: participants modified representations according to several characteristics.

- Encoding parameters: Two participants adjusted *encoding parameters*, in particular color and transparency. One participant had misunderstood the motion direction and had used categorical colors opposite to their intention (for speed up rather than slow down), which they then fixed. Another participant changed a color to one they felt fit better for the moving representations.
- Representation type: The use of `text` in full-motion mode dropped significantly. Three participants reported finding `graphics` to be more readable and easy to track than text under motion. Also, as some numerical data, including lap time, did not update per frame, designers felt the accuracy of text was not needed.
- Embedding location: Three participants changed the embedding locations of their designs. Two stated that they could not tell the swimmers' swimming direction from the still image. Thus, they put their designs in the wrong position relative to the swimmers. For example, they had planned to add national flags behind swimmers but put them in front instead by accident. One designer explained that they found it unnecessary to put the representations too close to swimmers as the motion trajectory of visualizations was exactly the same as the one of the swimmers. Thus, they could track and identify information even if there was a gap between visualizations and swimmers. Finally, one designer said that it simply seemed natural to modify the representation location to fit the video context better.
- Motion status: Three participants switched some representations' motion status. Two mentioned that their data did not change over time or based on the swimmers' position; thus, they did not have a good reason to let the representation move with swimmers. One participant commented that although `metadata` was related to the swimmer, situating it with the `swimming pool` was a better choice, as it would free the space around the swimmer for data that would update in real-time. One participant changed a static visualization to be in motion without giving a specific reason.

D. Results: Design Mode Preference

Here, we first report on participants' shifts in satisfaction between design modes as well as design mode preferences (Table III). We then, according to our second interview, explain in detail what led to the shifts in satisfaction and choice of design mode preference.

Shifts in Satisfaction & Design Mode Preference: Half 4/8 of our participants decreased their satisfaction after seeing how their static design looked on the video. 3/8 designers reported similar satisfaction as before. Only one designer was more satisfied with their static design. On the other hand, almost all designers 7/8 were more satisfied with their design produced in the full-motion mode. When asked, compared with the designs from the motion-limited mode (after seeing the rendering video), to what extent their design had been improved with the full-motion mode, 7/8 of the participants stated that their designs were improved by 1–4 points, out of 7-Likert points. One participant declined to respond. More telling is the overall preference of the design mode: All designers preferred to design with `full-motion mode`.

Reasons for Satisfaction Shift & Design Mode Preference: *Shifts in satisfaction:* In `motion-limited mode`, the four participants who decreased their satisfaction after seeing the played video mentioned the following reasons: (a) two did not recognize the correct moving direction and put visualizations in the wrong place; (b) one found that they had added too much text, which was hard to read while playing the video; (c) another designer complained that the motion effect did not match their expectation — the world record line was behind swimmers at the beginning of the race. This happened because the moving speed of the record line was calculated by dividing the total distance by the record time (as is done in the Olympics). Thus, the line moves at the average speed of the record keeper instead of their real-time speed, which can confuse audiences. The three participants who maintained their satisfaction commented that the played video largely matched their expectation, especially when seeing the motion effects matching with the moving entities in the video. Thus, they were not surprised or frustrated. The participant who largely increased their satisfaction said they were at first not satisfied with their designs before seeing the played video because the design method was totally different from what they were used to — sketching on a storyboard. They disliked the predefined visualizations and wanted to customize all visualizations by themselves. However, after seeing the video, they found the video "super cool" — all their designs had life and could move with swimmers or stay static on the swimming pool. They were so excited to see the rendered video because it was largely beyond their expectation. Thus, they gave the highest satisfaction, not for the design, but for the rendered video.

In `full-motion mode`, the seven participants increased their satisfaction because they had the opportunity to manipulate their designs in the same context in which their final design would be used. They mentioned that the full-motion mode gave them immediate visual feedback, which encouraged them to polish their designs and helped them to meet their designs expectations; much less mental effort was demanded of them, and participants appreciated that they no longer had to speculate about what designs would look like. The participant who did not rate their satisfaction explained that, without comparison, they did not know how to evaluate their satisfaction as they had never used a similar design method.

Design Mode Preference: We observed that designers played videos in the `full-motion mode` for two reasons: (1) to preview the designs they created on the video screenshots and (2) to change designs while the video was playing. The second case was the most common case.

`full-motion mode` was reported as the preferred design mode for the following benefits:

- Motion identification: Five participants said that the `full-motion mode` helped them to identify the correct moving direction, select which data items to add, let them choose better where to put things, check the conflicts between visualizations as well as overlap with swimmers, and get more feedback.
- Accessibility to dynamic data updates: Five participants reported that in the `full-motion mode`, they became aware of the updating frequency of a data item, which helped them to assign the movement status per item and check how the corresponding visualization would change. Participants were able to identify and de-prioritize data items that were updated less frequently than expected (*e.g.*, lap time) and switched them either to static representations or to simpler visuals.
- Flexibility of motion control: Four participants mentioned that with the flexibility of pausing, they could polish their design in detail and check design behavior at specific moments, such as turning.
- Instantaneous preview: Three participants said that as a designer, they should understand and be certain about their design before publishing it to the public. Thus, it was important to be able to have a quick design-reflect-redesign phase — previewing the effect of their design and making modifications.
- Context awareness and confidence: Three participants also much appreciated directly editing in the same context as the final product, which gave them confidence about their design as well as the final artistic product.
- Expectation match: Two participants stated that it would seem natural to do manipulations directly on the video if they were asked to design things for a motion context.

On the other hand, `motion-limited mode` was commented on in almost the inverse way. Participants found it hard to imagine the final product, and their design to be uncertain, unreliable, or confusing, they also misunderstood the swimming direction,and found it hard to avoid conflicts.

E. Results: Authoring Probe Improvement

During the second interview, participants also requested additional functions for *SwimFlow*. A set of features they wanted were, unsurprisingly, related to the more flexible specification of the visual encoding, akin to modern visualization authoring tools, which were too numerous to include in our probe. Some requests, however, were also related to the motion context and the situated visualization. Participants wanted some features akin to video editing tools that would allow them to select when representations would be visible. Others wanted visual effects like zoom-in/out or fade-in/out features to draw attention. Participants also voiced a wish for more data referents to which visualizations could be attached to, for example, to place a visualization relative to another visualization's position.

VI. Discussion

Overall, we found that the full-motion design context, and in particular the ability to see design changes reflected immediately with moving referents, was extremely helpful. The ability to play video and embedded visualizations together increased designers' satisfaction with their final designs. As such, it seems an important area for future work to develop prototyping and production tools for in-situ visualization in motion scenarios. Yet, several considerations arose from our work and the study results:

A. Impact of Context-Coupled Design on Vis Choices

When we started our work, we expected that the moving context (swimmers, background) would affect the design experience, as most authoring tools usually start with a blank canvas. We began our exploration on designing embedded visualizations in motion using traditional mechanisms that involved sketching on transparent overlays that allowed to move visualizations over a background image [8]–[10]. We observed that when designing on such a static frame, even when movement could be simulated, it was hard for designers to imagine the potential change of movement of the referents and the dynamics of the data attached to them. We observed similar problems in the motion-limited context when designers used our probe on a static video frame. Often, issues with visualization choices (color, transparency) and their placement were only noticed when designers could preview their designs, the data updated, and the corresponding entities moved together.

The importance of the playable video to design situated visualizations in motion is supported by the fact that all of the designers tweaked their visualizations, and more than half added/removed visualizations after seeing the video preview. Benefits of a video preview have been found in other tasks and contexts, such as turning visualizations on/off by audiences [16], [38]. Here we are able to explain where benefits come from. We saw, for example, that the availability of a full-motion design mode influenced the selection of representations. Some designers were reluctant to add multiple moving visualizations when designing with incomplete contextual motion, and often added text as the simplest representation. Nevertheless, when they actually saw their completed design in motion with the referent, they felt that the visuals were more subtle than expected and that they could add more moving components. Designers also found that graphics were easier to track than text when they were moving. Thus, context-coupled design with complete motion effects impacted not just simple visual design choices (such as color or position) but also data representations.

Arguably, the most surprising were visualization changes due to the data values themselves. While playing their designs in the video, several designers noticed that some data elements they considered important and were often attached to the swimmer, were, in fact, not updating regularly or not as expected. Thus, designers decided it did not make sense to give them prominent positions, such as the regions around the swimmer, and either removed them or reverted to static visualizations. We plan to investigate in future research if these findings apply in other moving contexts, for example, in sports or games where trajectories are not necessarily as linear as in swimming.

As a disclaimer: some of our findings related to design choices may have been impacted by our study design. We did not show the target video to participants before the task started. Instead, we showed a short Olympic video to give context. Not having the final video meant that participants did not need to try and memorize possible events and track performance when they were designing on the static frame. We do not know how this choice impacted final designs; for example, if designers had memorized key events, they might have placed visualizations differently and they may have changed their preferences.

While we found advantages to the full-motion design mode, future tools inspired by our findings will not necessarily replace traditional methods such as sketching for prototyping. For example, people could sketch a set of first ideas using their preferred method (paper-based designs, storyboards, video presentations, cardboard mockups, etc.) and then switch to a design tool to try them out and refine them. We can also consider intermediate prototyping approaches, adapting sketching methodologies to a motion context; for example sketch-based authoring tools that allow designers to sketch visualizations on videos directly. Existing sketch-based authoring tools (e.g., [39], [51]), can indeed combine sketching with data-binding. However, they need to be adapted to preview videos and to provide object and trajectory detection in order to help designers attach their sketched visualizations to data referents.

Summary: The study with our technology probe is a first exploration focused on how to *design* embedded visualizations in a motion context and, in particular, with visualizations that are meant to move with their referents. The most significant finding from the study is the importance for designers to see the motion of referents together with visualizations but, very importantly, also to see how the visualized data updated with movement. While our probe only targeted the swimming context, we hypothesize that the importance of the full-motion mode will also hold for design environments that target other types of sports, in particular, if these other sports include very complex and dynamic types of movement. Yet, with other types of motion and contextual factors in other sports, a future tool might need dedicated features to address these complexities. The combination of sketching and full-motion contexts is an interesting future avenue for research on the design process for embedded visualizations in motion.

B. Tool Complexity

Our study was based on a technology probe which, by design, is not a complete system but instead is meant to help envision what a *complete* visualization authoring tool for a motion context would look like. Specifically, participants expressed the need for visualization authoring features found in complete authoring tools such as Charticulator [35], Lyra [33], or Data Illustrator [32]. But they also requested features from full video editing tools, such as zoom in/out, fade in/out, control of timings, etc. In addition to these two sets of features, we also need to provide ways for designers to embed and anchor their visualizations to particular moving objects (as we do in our technology probe), assuming such objects are pre-identified. If they are not identified, designers also need to define trajectories for their moving visualizations. There may even need to be features to define the updating of dynamic data (like the current swimmer speed) as the video is playing: how often to update which data, how to highlight changes, etc.

Furthermore, we need to further study the impact of these features for the visualization designer and reader. Swimming mainly contains linear trajectories, but irregular trajectories are more common in other scenarios, such as ball sports and video games. Irregular movement can lead to overlap between entities. How to support designers to overcome such overlap between embedded visualizations is still a challenge. The moving speed is another aspect that might affect design decisions. Previous work from Yao et al. [1] showed that not only irregular trajectories, but also higher speeds will lead to more errors in visualization reading. How to design readable visualization in a motion context, especially with variation speeds, is also an open question. Besides, the movement direction might also impact the visualization design, especially for arrow-like icons. For example, in swimming, turnings might happen non-synchronously between swimmers and, as such, can lead to two visualizations, such as arrows, facing opposite directions. Similar situations where various movement directions are involved, also exist in other sports, such as team sports with many moving teammates (soccer, basketball, etc.). How to help audiences avoid direction misunderstanding still needs further tests and explorations.

A large number of functionalities are desirable to support the design of visualizations in motion. However, are all features essential? What are, in the end, the key features for such authoring tools to avoid overloaded interfaces and feature fatigue? Should one single authoring tool attempt to cover the entire design process, or is a multi-step design workflow more suitable (defining the video/motion-related features in one tool and refining the visualization design details in another)? These all remain open questions.

An added problem in the support of a design workflow is the requirement of an underlying video with associated tracking data. In our case, the tracking data was prepared semi-automatically in that some data was extracted automatically, and other data had to be annotated by hand. For example, we had no computer vision algorithm that could detect a swimmer's single stroke. Future tools will similarly have to rely on the availability of data or include possible manual annotation features. Systems such as iBall [20], for example, put a lot of effort into the development of a computer vision pipeline to track players and game information for embedding visualizations. When a future authoring tool needs to support multiple camera shots and movement, having tracking data synchronized and registered from all cameras becomes important. The tool complexity would further grow with the required need to detect or define perspective distortions of the scene so that visualizations can be embedded correctly.

Summary: We developed *SwimFlow* as a technology probe, and as such, it worked very well to elicit features of future technology to support design processes around embedded visualizations in sports videos. A big problem for fully-fledged tools will be the potential necessity to be visualization authoring tools, with video editing features, data tracking and editing possibilities, and new features for controlling how visualizations will be embedded in relation to the data referents.

C. Guidance on Effectiveness

In our work, we used a user-centered design process (survey, workshop, testing with designers) to prioritize the data and the visual representations to include in our technology probe. There are many more possibilities in terms of data but also visual representations to support. These can be, of course, added if our probe was turned into a future prototype. One potential future feature our technology probe did not concretely surface is the potential need to provide designers with guidance on how best to design embedded visualizations in motion. For static visualization design, the community has a large number of design recommendations for how to best represent data for different tasks and contexts, and even recommender tools [52]–[54] and catalogs [55], [56]. Yet, it is unclear if our current understanding of how effective visualizations are, still applies to motion contexts. As such, a future recommender system in this context would be based on little empirical evidence. Recent work [1] has started investigating the impact of motion on reading visualizations, but clearly, more work is needed to be able to provide concrete design recommendations for effective visualizations in motion.

Summary: In our study, we compared full-motion to a motion-limited design mode. The motion-limited mode is fundamentally the same as offering designers of embedded visualizations a screenshot as a design prop to create visualizations on. This would be a low-cost design alternative that does not necessarily require videos or all tracking data - but could simply use authoring and embedding features. Our study explored to what extent having the full-motion mode would be useful and why. Specifically, we found that the full-motion mode led to more diverse designs, less mental effort, a faster design process, reduced uncertainty / more confidence about the final design, and more satisfaction. The full-motion mode allowed the preview of dynamically updating data and gauging the potential visual overload of multiple embedded visualizations in motion. Interestingly, participants also changed representation types when they saw visualizations moving. Text became much less commonly used, indicating the difficulty of imagining what representations might look like when in motion. Finally, previewing motion was important for several participants who had misunderstood movement direction.

VII. Conclusions

We investigated how to design visualizations that may be attached to moving contexts, such as swimmers in a race. This type of embedding is becoming increasingly prevalent in sports videos to communicate metadata about athletes' performance and race/game metadata. To explore the design context for visualizations in motion, we focused on swimming. We first identified the data swimming enthusiasts wanted to see in videos through an online survey. We then used this data in a design workshop to identify appropriate visualization designs that can be embedded in swimming videos and to see the first difficulties with low-fidelity prototyping. These designs were then used to inform the creation of *SwimFlow* a technology probe we built to understand how the motion contexts affect the design process of visualization designers.

Through this process, we provide insights about what data and visualizations are a good fit for the swimming context. More importantly, we identified how the motion context can greatly affect the visualization design process and choices. We discussed the more general challenges and open questions in adapting existing prototyping and visualization authoring practices for video contexts and pointed to a vast space to explore how to design situated visualizations in motion.

We invite more visualization researchers, as well as those who have an interest in motion features and context, to join in this emerging research direction. There are many factors to explore related to the perception of embedded visualizations in motion that may eventually lead to guidelines: various application scenarios in and outside of sports, different target audiences and their focus, the data, and, of course, the visualizations in motion themselves, among others.

Images/graphs/plots/tables/data license/copyright

Acknowledgment

This work was partly supported by the Agence Nationale de la Recherche (ANR), grant number ANR-19-CE33-0012 and ANR-19-STHP-0004. The ethics approval number of the work is CER-Paris-Saclay-2020-074/amendement1.

Supplemental Material Pointers

The pre-registrations for our (a) survey and (b) user study can be found at (a) `https://osf.io/qdbhg` and (b) `https://osf.io/fw7gj`. Respectively. We also share our (a) survey structure, results, and analysis scripts; (b) sketches collected from the workshop; (c) user study results, including created designs in both static and video mode in video format; (d) recorded swimming videos under the authorization, as well as codes to create and combine videos; (e) source codes and demonstration video of *SwimFlow*; and other materials (appendix, ethics approval...) at `https://osf.io/nxyr4/`. *SwimFlow* can be accessed at `https://motion.isenberg.cc/Swimming/index`.

References

[1] L. Yao, A. Bezerianos, R. Vuillemot, and P. Isenberg, "Visualization in motion: A research agenda and two evaluations," *IEEE Trans. Vis. Comput. Gragh*, vol. 28, no. 10, pp. 3546–3562, 2022.

[2] FootoVision, https://www.footovision.com/ Accessed 08 2021.

[3] Tokyo 2020 Mixed-Gender 4×100m Medley Final, https://www.youtube.com/watch?v=Ye2wK3ClvO4&t=505s Accessed 02 2023.

[4] Tokyo 2020 Women's 100m Freestyle Final, https://www.youtube.com/watch?v=EPPrQFF_TTY&t=221s Accessed 02 2023.

[5] Tokyo 2020 Men's 100m Butterfly Final, https://www.youtube.com/watch?v=Mu7ZmJduQPI&t=287s Accessed 02 2023.

[6] SportBuzzBusiness, https://www.youtube.com/watch?v=fVB1QeiDy6Q&t=10s Accessed 02 2023.

[7] T. Lin, Z. Chen, Y. Yang, D. Chiappalupi, J. Beyer, and H. Pfister, "The quest for Omnioculars: Embedded visualization for augmenting basketball game viewing experiences," *IEEE Trans. Vis. Comput. Gragh*, vol. 29, no. 1, pp. 962–971, 2023.

[8] H. Park and S. McKilligan, *A Systematic Literature Review for Human-Computer Interaction and Design Thinking Process Integration*. Cham, Switzerland: Springer, 06 2018, pp. 725–740.

[9] Y. Rogers, H. Sharp, and P. Jennifer, *Interaction design: Beyond Human-Computer Interaction*, 6th ed. Wiley, 04 2023.

[10] A. Dix, *Human-computer interaction*, 3rd ed. Pearson, 2003.

[11] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, N. Roussel, and B. Eiderbäck, "Technology probes: Inspiring design for and with families," in *Proc. Conf. Hum. Factors Comput. Syst.*, 2003, pp. 17—24.

[12] W. Willett, Y. Jansen, and P. Dragicevic, "Embedded data representations," *IEEE Trans. Vis. Comput. Gragh*, vol. 23, no. 1, pp. 461–470, 2016.

[13] N. Bressa, K. Wannamaker, H. Korsgaard, W. Willett, and J. Vermeulen, "Sketching and ideation activities for situated visualization design," in *Proc. Conf. Design. Interact. Syst.* ACM, 06 2019, pp. 173–185.

[14] T. Lin, R. Singh, Y. Yang, J. Beyer, and H. Pfister, "SportsXR: Immersive Analytics in Sports," in *Proc. Immersive Analytics Workshop*, 04 2020.

[15] T. Tang, J. Tang, J. Hong, L. Yu, P. Ren, and Y. Wu, "Design guidelines for augmenting short-form videos using animated data visualizations," *J. of Vis.*, vol. 23, no. 4, pp. 707–720, Aug. 2020.

[16] Z. Chen, S. Ye, X. Chu, H. Xia, H. Zhang, H. Qu, and Y. Wu, "Augmenting sports videos with VisCommentator," *IEEE Trans. Vis. Comput. Gragh*, vol. 28, no. 1, pp. 824–834, 2022.

[17] M. Stein, H. Janetzko, A. Lamprecht, T. Breitkreutz, P. Zimmermann, B. Goldlücke, T. Schreck, G. Andrienko, M. Grossniklaus, and D. A. Keim, "Bring it to the pitch: Combining video and movement data to enhance team sport analysis," *IEEE Trans. Vis. Comput. Gragh*, vol. 24, no. 1, pp. 13–22, 01 2018.

[18] SportsDynamics, https://sportsdynamics.eu/ Accessed 08 2021.

[19] F. Bucchieri, "Situated visualization in motion for video games," Master's thesis, Université Paris-Saclay and Politechnico Milano, France, Italy, 2022, http://hdl.handle.net/10589/195726.

[20] Z. Chen, Q. Yang, J. Shan, T. Lin, J. Beyer, H. Xia, and H. Pfister, "IBall: Augmenting Basketball Videos with Gaze-Moderated Embedded Visualizations," in *Proc. Conf. Hum. Factors Comput. Syst.*, no. 841, 04 2023, pp. 1–18.

[21] C. Perin, R. Vuillemot, C. D. Stolper, J. T. Stasko, J. Wood, and S. Carpendale, "State of the Art of Sports Data Visualization," *Comput. Graph. Forum*, vol. 37, no. 3, pp. 663–686, 2018.

[22] H. Oagaz, B. Schoun, and M. Choi, "Performance improvement and skill transfer in table tennis through training in virtual reality," *IEEE Trans. Vis. Comput. Gragh*, vol. 28, no. 12, pp. 4332–4343, 2022.

[23] J. Wang, J. Wu, A. Cao, Z. Zhou, H. Zhang, and Y. Wu, "Tac-Miner: Visual tactic mining for multiple table tennis matches," *IEEE Trans. Vis. Comput. Gragh*, vol. 27, no. 6, pp. 2770–2782, 2021.

[24] E. Wu, M. Piekenbrock, T. Nakumura, and H. Koike, "Spinpong - virtual reality table tennis skill acquisition using visual, haptic and temporal cues," *IEEE Trans. Vis. Comput. Gragh*, vol. 27, no. 5, pp. 2566–2576, 2021.

[25] S. Chen, B. Hu, Y. Gao, Z. Liao, J. Li, and A. Hao, "Lower limb balance rehabilitation of post-stroke patients using an evaluating and training combined augmented reality system," in *Adjunct Proc. Symp. Mixed and Aug. Reality*, 2020, pp. 217–218.

[26] C. Perin, R. Vuillemot, and J. D. Feket, "Real-tima crowdsourcing of detailed soccer data," in *Proc. Sports Data Vis. Workshop*, 10 2013.

[27] T. Lin, R. Singh, Y. Yang, C. Nobre, J. Beyer, M. A. Smith, and H. Pfister, "Towards an understanding of situated AR visualization for basketball free-throw training," in *Proc. Conf. Hum. Factors Comput. Syst.*, 05 2021, pp. 461:1–461:13.

[28] L. Bai, C. Efstratiou, and C. S. Ang, "weSport: Utilising wrist-band sensing to detect player activities in basketball games," in *Proc. Conf. Pervasive Comput. and Comm. Workshop*. IEEE, 2016, pp. 1–6.

[29] G. Pingali, A. Opalach, Y. Jean, and I. Carlbom, "Visualization of sports using motion trajectories: providing insights into performance, style, and strategy," in *Info. Vis.* IEEE, 2001, pp. 75–544.

[30] S. Ye, Z. Chen, X. Chu, Y. Wang, S. Fu, L. Shen, K. Zhou, and Y. Wu, "ShuttleSpace: Exploring and analyzing movement trajectory in immersive visualization," *IEEE Trans. Vis. Comput. Gragh*, vol. 27, no. 2, pp. 860–869, 2021.

[31] T. Polk, J. Yang, Y. Hu, and Y. Zhao, "TenniVis: Visualization for tennis match analysis," *IEEE Trans. Vis. Comput. Gragh*, vol. 20, no. 12, pp. 2339–2348, 12 2014.

[32] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu, "Critical reflections on visualization authoring systems," *IEEE Trans. Vis. Comput. Gragh*, vol. 26, no. 1, pp. 461–471, 2020.

[33] A. Satyanarayan and J. Heer, "Lyra: An interactive visualization design environment," *Comput. Graph. Forum*, vol. 33, no. 3, pp. 351–360, 2014.

[34] Z. Liu, J. Thompson, A. Wilson, M. Dontcheva, J. Delorey, S. Grigg, B. Kerr, and J. Stasko, "Data Illustrator: Augmenting vector design tools with lazy data binding for expressive visualization authoring," in *Proc. Conf. Hum. Factors Comput. Syst.* ACM, 2018, pp. 123:1–123:13.

[35] D. Ren, B. Lee, and M. Brehmer, "Charticulator: Interactive construction of bespoke chart layouts," *IEEE Trans. Vis. Comput. Gragh*, vol. 25, no. 1, pp. 789–799, 2019.

[36] Y. Shi, P. Liu, S. Chen, M. Sun, and N. Cao, "Supporting expressive and faithful pictorial visualization design with visual style transfer," *IEEE Trans. Vis. Comput. Gragh*, vol. 29, no. 1, pp. 236–246, 2023.

[37] J. E. Zhang, N. Sultanum, A. Bezerianos, and F. Chevalier, "Dataquilt: Extracting visual elements from images to craft pictorial visualizations," in *Proc. Conf. Hum. Factors Comput. Syst.* ACM, 2020, pp. 45:1—-45:13.

[38] Z. Chen, Q. Yang, X. Xie, J. Beyer, H. Xia, Y. Wu, and H. Pfister, "Sporthesia: Augmenting sports videos using natural language," *IEEE Trans. Vis. Comput. Gragh*, vol. 29, no. 1, pp. 918–928, 10 2023.

[39] R. Suzuki, R. H. Kazi, L.-y. Wei, S. DiVerdi, W. Li, and D. Leithinger, "RealitySketch: Embedding responsive graphics and visualizations in ar through dynamic sketching," in *Proc. Symp. User Interface Software and Tech.* ACM, 2020, pp. 166—-181.

[40] Z. Chen, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu, "MARVisT: Authoring glyph-based visualization in mobile augmented reality," *IEEE Trans. Vis. Comput. Gragh*, vol. 26, no. 8, pp. 2645–2658, 2020.

[41] D. Shi, F. Sun, X. Xu, X. Lan, D. Gotz, and N. Cao, "Autoclips: An automatic approach to video generation from data facts," *Comput. Graph. Forum*, vol. 40, no. 3, pp. 495–505, 2021.

[42] T. Tang, J. Tang, J. Lai, L. Ying, Y. Wu, L. Yu, and P. Ren, "Smartshots: An optimization approach for generating videos with data visualizations embedded," *ACM Trans. Interact. Intell. Syst.*, vol. 12, no. 1, Mar. 2022.

[43] LimeSurvey, https://sondages.inria.fr/ Accessed 03 2023.

[44] L. Yao, A. Bezerianos, R. Vuillemot, and P. Isenberg, "Situated Visualization in Motion for Swimming," in *Journée Visu 2022*, Bordeaux, France, Jun. 2022. [Online]. Available: https://hal.inria.fr/hal-03700406

[45] J. C. Roberts, C. Headleand, and P. D. Ritsos, "Sketching designs using the five design-sheet methodology," *IEEE Trans. Vis. Comput. Gragh*, vol. 22, no. 1, pp. 419–428, 2016.

[46] Olympics Aquatics, "Welcome to the women's 200m butterfly final at the Olympics Tokyo 2020." https://www.youtube.com/watch?v=WcuGsLbnHqY Accessed 03 2022.

[47] B. Buxton, *Sketching User Experiences: Getting the Design Right and the Right Design*. Amsterdam: Morgan Kaufmann, 2011.

[48] FFN, "French national swimming federation," Visited 06 2023. [Online]. Available: https://www.ffnatation.fr/

[49] labelImg, "Welcome to labelImg." https://github.com/HumanSignal/labelImg Accessed 07 2023.

[50] SwimFlow, "The example design made and shared by swimflow," Created 01 2023. [Online]. Available: https://motion.isenberg.cc/Swimming/sharing?822634378679532_v68

[51] H. Xia, N. Henry Riche, F. Chevalier, B. De Araujo, and D. Wigdor, "DataInk: Direct and creative data-oriented drawing," in *Proc. Conf. Hum. Factors Comput. Syst.* ACM, 2018, p. 223:1–223:13.

[52] K. Hu, M. A. Bakker, S. Li, T. Kraska, and C. Hidalgo, "VizML: A machine learning approach to visualization recommendation," in *Proc. Conf. Hum. Factors Comput. Syst.* ACM, 2019, p. 128:1–128:12.

[53] V. Dibia and Ç. Demiralp, "Data2Vis: Automatic generation of data visualizations using sequence-to-sequence recurrent neural networks," *IEEE Comput. Graph. and App.*, vol. 39, no. 5, pp. 33—46, 2019.

[54] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer, "Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco," *IEEE Trans. Vis. Comput. Gragh*, vol. 25, no. 1, pp. 438–448, 2019.

[55] Y. Holtz and C. Healy, "From data to vis," Website: http://www.data-to-viz.com/, Accessed 2022-03-25. [Online]. Available: http://www.data-to-viz.com/

[56] S. Ribecca, "The data visualisation catalogue," Accessed 2022-03-25. [Online]. Available: http://datavizcatalogue.com/

Lijie Yao received her B.Eng. degree from Xidian University in 2018, her diplôme d'ingénieur (equivalent to M.Eng degree) from École polytechnique universitaire de Lille in 2019, and her M.S. degree from Université Paris-Saclay in 2020. She recently defended her Ph.D. thesis at AVIZ, Inria and Univrsité Paris-Saclay. Her main research areas are information visualization and visual analytics, with a focus on visualization in motion. Find more information about her & contact her at lijieyao.com.

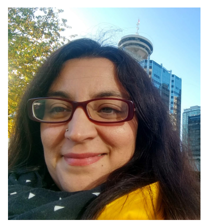

Anastasia Bezerianos is a full Professor at Université Paris-Saclay, in the Polytech engineering school. She is part of the LISN research lab, as well as a member of the ILDA Inria team. She works at the intersection of human-computer interaction and information visualization. Contact her at anastasia.bezerianos@universite-paris-saclay.fr

Romain Vuillemot is an Assistant Professor of Computer Science at École Centrale de Lyon, Université de Lyon, France, in the LIRIS laboratory. His research interests include information visualization and visual analytics, for application domains such as sport, traffic data, and machine learning explicability. Contact him at romain.vuillemot@ec-lyon.fr

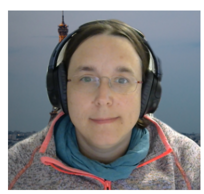

Petra Isenberg is a research scientist at Inria, France in the Aviz team. Her main research areas are visualization and visual analytics. She is interested in exploring how people can most effectively work when analyzing large and complex data sets—often on novel display technology. Contact her at petra.isenberg@inria.fr.