This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Saloheimo, Tuure; Kaos, Maximus; Fricker, Pia; Hämäläinen, Perttu

Automatic Recognition of Playful Physical Activity Opportunities of the Urban Environment

Published in:
Academic Mindtrek 2021

DOI:
[10.1145/3464327.3464369](https://doi.org/10.1145/3464327.3464369)

Published: 01/06/2021

Document Version
Peer reviewed version

Please cite the original version:

Saloheimo, T., Kaos, M., Fricker, P., & Hämäläinen, P. (2021). Automatic Recognition of Playful Physical Activity Opportunities of the Urban Environment. In *Academic Mindtrek 2021: Proceedings of the 24th International Conference on Academic Mindtrek* (pp. 49-59). ACM. <https://doi.org/10.1145/3464327.3464369>

Automatic Recognition of Playful Physical Activity Opportunities of the Urban Environment

Tuure Saloheimo

Aalto University
Finland

tuure.saloheimo@gmail.com

Pia Fricker

Aalto University
Finland

pia.fricker@aalto.fi

Maximus D. Kaos

University of Southern Denmark
Denmark

mkaos@health.sdu.dk

Perttu Hämäläinen

Aalto University
Finland

perttu.hamalainen@aalto.fi

Figure 1: Our neural network detects parkour spots from street level photographs. The figure shows example visualizations of the results. Left: map view with dataset images as black and green dots, green ones containing detected spots. The image closest to the mouse pointer is shown to allow quick browsing of the data. Middle: mouse click expands the image. Right: visualizing the top-scoring detections.

ABSTRACT

We investigate deep neural networks in recognizing playful physical activity opportunities of the urban environment. Using transfer learning with a pre-trained Faster R-CNN network, we are able to train a parkour training spot detector with only a few thousand street level photographs. We utilize a simple and efficient annotation scheme that only required a few days of annotation work by parkour hobbyists, and should be easily applicable in other contexts, e.g. skateboarding. The technology is tested through parkour spot exploration and visualization experiments. To inform and motivate the technology development, we also conducted an interview study about what makes an interesting parkour spot and how parkour hobbyists find spots. Our work should be valuable for researchers and practitioners of fields like urban design and exercise video

games, e.g., by providing data for a location-based game akin to Pokémon Go, but with parkour-themed gameplay and challenges.

CCS CONCEPTS

• **Human-centered computing** → **Geographic visualization**; *Empirical studies in HCI*; • **Computing methodologies** → *Scene understanding*.

KEYWORDS

machine learning, parkour, transfer learning, computer vision, urban design, playable cities

ACM Reference Format:

Tuure Saloheimo, Maximus D. Kaos, Pia Fricker, and Perttu Hämäläinen. 2021. Automatic Recognition of Playful Physical Activity Opportunities of the Urban Environment. In *Academic Mindtrek 2021 (Mindtrek '21)*, June 1–3, 2021, Tampere/Virtual, Finland. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3464327.3464369>

1 INTRODUCTION

To help people be physically active, the availability and accessibility of exercise is of high importance [9]. From this perspective, urban sports like parkour, skateboarding, and street workout are highly relevant, as they can be practiced almost anywhere with

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Mindtrek '21, June 1–3, 2021, Tampere/Virtual, Finland

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8514-5/21/06...\$15.00

<https://doi.org/10.1145/3464327.3464369>

Figure 2: An example of parkour practice suitable for both beginners and more experienced traceurs. Captured at an open parkour jam in Helsinki, demonstrated by <https://www.instagram.com/taavetsu>.

only modest or no cost. However, it can be hard to find training spots/areas that are both interesting and conveniently located. As lack of time is one of the perceived barriers to physical activity [8, 49], it would be desirable to minimize the time spent traveling to and from exercise. Greater distance to exercise has also been empirically linked to lower physical activity [42].

In this paper, we tackle the exercise opportunity discovery problem in the case of parkour. Our aim is to augment the sport with novel tools that allow locating interesting parkour spots. Although common urban geometry like railings, lamp posts, and staircases all provide some parkour affordances, the most interesting and enjoyable parkour spots are rare. Our focus on parkour is motivated by the domain expertise of one of the authors, and an active local parkour community that organized weekly "parkour jams" open to anyone, allowing us to observe and interview parkour hobbyists. Parkour is also a highly accessible form of physical activity that can adapt to different levels of skill and fitness. Extreme videos where parkour athletes leap and vault between rooftops are not representative of an average practice session, in the same way as Olympic gymnastics is not representative of gymnastics practiced as a hobby. Figure 2 shows an example of typical beginner or intermediate parkour practice, which happens safely near the ground but nevertheless develops strength, coordination, and spatial understanding.

This paper investigates the following research questions:

- What kind of training spots do parkour hobbyists or *traceurs* find particularly interesting or enjoyable?
- How do traceurs discover training spots?
- To facilitate the discovery of interesting spots, can one combine computer vision and large-scale image databases like Flickr or Google Street View to automatically recognize and visualize parkour spots?
- More specifically, can one find a technological approach that only requires a minimal amount of labeled training data? This matters because modern computer vision systems are typically trained with thousands or even millions of images labeled through crowdsourcing services such as Amazon Mechanical Turk. For example, the widely used ImageNet dataset [40, 47] provides over a million images with bounding box object annotations. However, if the labeling requires expert knowledge such as an understanding of parkour affordances, large-scale crowdsourcing is not feasible.

We answer the first two research questions through interviewing and observing parkour hobbyists. The central findings are that interesting parkour spots are akin to small *playgrounds*, offering high variety of shapes and affordances, and to find good spots, it is not enough to simply recognize specific types of geometry. Traceurs find spots through word-of-mouth, social media, exploring the city, and also through existing community-maintained online maps. However, such maps are not available for many locations, which motivates the development of automatic discovery and mapping tools.

Regarding the last two questions, we propose and evaluate a deep neural network approach for automatically detecting interesting parkour spots in street level images. We train the network using Google Street View images and demonstrate that it generalizes to test data we captured ourselves. We utilize a simple and fast image labeling scheme together with a transfer learning approach, only requiring a few thousand images and a few days of labeling work by domain experts (parkour hobbyists in our case). Source code is available at <https://github.com/ThetaNord/parkour-detection>.

In summary, we make the following contributions:

- We present a novel application of machine learning, in the form of a system for detecting and visualizing interesting parkour spots from street level imagery. Our work demonstrates how modern machine learning tools may be leveraged to open one's eyes to new playful properties of the urban environment.
- We advance the understanding of what makes an interesting parkour spot and how traceurs find them.

Our findings and the data generated by our system could enable novel location-based games akin to Pokémon Go [32], but with parkour-themed gameplay and challenges. Our technological approach should be easily applicable to other contexts such as skateboarding. Beyond games and play, automatic collection of geographic exercise opportunity information should be useful for urban design and research that investigates how the built environment affects physical activity. For example, the study of Estabrooks et al. [9] operationalized exercise availability as the presence of exercise resources within a given neighborhood; our approach allows collecting such data about a wider palette of physical activities. Previous analyses have focused on traditional measures such as neighborhood walkability and the presence of parks and recreation facilities [4].

2 BACKGROUND AND RELATED WORK

2.1 Understanding Parkour

Parkour is an activity revolving around moving through and around obstacles, often in urban spaces, with speed and efficiency. Parkour practitioners — so-called "traceurs" — chain together movements (i.e., leaping, vaulting, balancing over obstacles) to achieve a "flow path" through the environment [35]. Despite its emphasis on efficiency, parkour has been described as a "form of unscripted creative play" that reinterprets the city as "a terrain of playful possibility" ([5], p. 393), as well as "a declaration of the creative ludic potential and the playing spirit of mankind" ([37], p. 21). For the traceurs interviewed by Ameel & Tani [2], parkour seemed to be a way of continuing to use public space in a childlike manner, and Leone [22]

goes as far as describing parkour as a desire to turn the entire city into an entertainment park. Thibault [50] elaborates this through a threefold playful characterization: playing *in* the city, playing *with* the city, and also *playing the city* by escaping and opposing its logic.

While some texts like Leone [22] and Thibault [50] view parkour through a polemic and political lens, e.g., as resistance to top-down urban planning, others note a trend towards parkour becoming mainstream and institutionalized with national associations, paid classes, and instructor training [2]. As an example of the latter, the last author of this paper and his children have taken parent-child parkour classes at a local parkour school, and view parkour simply as a form of exercise that provides a particularly interesting combination of variety, creativity, physical intensity, and everyday practical applicability.

Related to the parkour spot discovery problem solved in this paper, traceurs are willing to go far out of their way to frequent spaces that are known to be well-suited to their parkour practice [1]. On the other hand, parkour novices and non-practitioners do not perceive urban spaces as traceurs do; through practicing parkour, traceurs learn to re-interpret and perceive which environments are most safe, interesting and even aesthetic [1]. This notion is sometimes described as gaining "parkour vision" [43] or "parkour eyes" [1], which reveals new dimensions in one's surroundings [1], turns limitations into opportunities [2], and allows more playful options to emerge [43].

In light of the above, teaching computers "parkour vision" seems a worthwhile and interesting challenge, furthering and disseminating the understanding of how traceurs perceive their surroundings. This paper takes the first steps towards this goal, and our interview study about parkour spots also adds detail to previous work on what traceurs find interesting [1, 2, 37, 43].

2.2 HCI of Physical Exercise and Play

Our work is about developing novel technology tools for playful physical activity. In the HCI literature, there exists a large body of related work on experimental exercise systems and movement-based games [16–18, 28, 30]. Beyond systems and case studies, the field has been pushed forward in the form of conceptual frameworks for sport and exercise design [12, 25, 29], and through increasing the understanding of movement-based game user experience [7, 15] and embodied playful activity beyond games [24, 52].

A portion of the HCI literature expands the discussion to non-digital physical play [12, 26], or physical exercise with relatively subtle technological augmentation, such as electrically assisted bicycling [3] or designing services for communities such as traceurs [51]. In this vein, we also augment the exercise experience with technology, focusing on the pre-exercise phase of deciding what to practice and where. Our results could be integrated with an existing parkour community service or developed into a dedicated app akin to the ones used by climbers to find climbing routes¹.

2.3 Urban Design and Play

Parkour has been framed as urban play, exemplifying how playful activity can contribute to the image of a city physically, socially

and culturally [37]. Various urban gamification solutions demonstrate how game-like experiences can kindle new urban joy [27]. Urban games and gamification case studies are too numerous to review here; instead, we refer the reader to the more general articles by Thibault [50] and Nijholt [33, 34]. Using the terminology of Thibault's typology of urban gamification [50], we regard machine learning as a tool for supporting bottom-up "urban writing": We model traceurs' playful perceptions and interpretations of the urban fabric, and disseminate them through digital means. Considering Nijholt's discussion of playable cities and various sensors and actuators for implementing playful digital smartness [33, 34], our work repurposes street level imagery as the sensing technology, and reveals emergent, non-designed potential for Third Places, where people can meet in a playful mood outside home or work, in our case for parkour jams and practice.

The Sustainable Development Goals (SDGs) of the United Nations [31] are challenging architects and urban planners to extend their knowledge and design practices, e.g., in relation to the topics of ensuring healthy lives and promoting well-being for everyone (SDG 3) and making cities and human settlements inclusive, safe, resilient and sustainable (SDG 11). At the same time, enabling people to understand the potentials of built and unbuilt environments leads to a new form of user engagement [45]. This calls for new methods for spatial and social analysis of complex urban environments, e.g., the development of novel Geographic Information Systems (GIS) [48]. From this perspective, our work provides a novel case study of using technology to unveil the play and physical activity potential of the built environment, extending the consideration of physical activity affordances like walkability [4] to the domain of "parkourability". Our system provides a new type of data that could be used by various GIS tools and models.

2.4 Analyzing Urban Imagery

With modern deep convolutional neural networks, basic image classification is fairly straightforward [11, 21]. Recognizing individual objects from images is somewhat more complicated, but nevertheless accessible thanks to open source packages like the Detectron [10]. We use a pre-trained Detectron network fine-tuned with our custom data; hence, our technical contribution is not on deep learning methods as such but in demonstrating that our novel use case is feasible with our chosen network architecture, dataset, and annotation approach.

There exists previous work that has used street level photographs as an information source for geographical information analysis, with either automated or crowdsourced data, to identify accessibility problems [13], landmarks for pedestrian navigation [20], curb ramps [14], and scenic driving routes [39], but we know of no previous work on recognizing parkourability or other physical activity opportunities.

3 INTERESTING PARKOUR SPOTS

Although the existing literature provides rich ethnographic accounts of how parkour practitioners experience their surroundings [1, 2, 37, 43], more specific information is lacking about what kinds of training spots are particularly interesting or enjoyable. To address this knowledge gap, we conducted a brief interview study

¹e.g., <https://27crag.com/>

with parkour hobbyists in Helsinki, Finland. The data informed our choice of machine learning and data annotation approach and provides validation for the need and usefulness of parkour spot maps.

3.1 Procedure and Method

We interviewed the participants about their favorite spots, querying for the following:

- What makes your favorite spots interesting or enjoyable?
- How do you find parkour spots (friends, spot maps, exploring the city, other)?
- Why do you like parkour?

Three traceurs answered the questions online, recruited through posting the questions to a local parkour Facebook group. Additionally, one of the authors — himself a parkour hobbyist — recruited participants via attending two parkour jams, i.e., open free-form practice sessions. The interview structure followed the online form with the above questions, and respondents received a choice of sports drink. In total, we received 9 responses. The mean age of the respondents was 26 years and they had 6 years of parkour experience, on average. Online recruitment turned out particularly difficult, and it was easier to approach people in person in the parkour jams. However, the number of active jammers turned out to be fairly small. In the first jam, we interviewed all 5 participants. The second jam only had 3 participants, out of which only one had not already answered the questionnaire. Nevertheless, even the 9 responses reveal clear patterns, as discussed below, and recruiting at parkour jams allowed us to observe traceurs in action and collect image material.

3.2 Results

We conducted an open coding [41] of participants' responses with regards to the questions. Note that overall responses were rather short, as it was impossible to inquire further via the online survey, and also because participants at the parkour jams wished to spend most of their time practicing.

3.2.1 *Favorite Parkour Spot Qualities* . A clear result was that while many urban shapes like rails, stairs, and lamp posts all provide some parkour opportunities, traceurs prefer spots with a high degree of variety. This was indicated by all answers, e.g., *"That the spot provides opportunities for many types of movements", "Variety, being at least partly usable in summer, autumn and spring, works for beginners, reachable by public transport and by bike. Can be used to practice balancing, skipping, jumping, climbing."*, *"Versatility, different shapes, different height levels"*, *"Rails, walls, wall run places, obstacles in different angles."* The importance of variety is in agreement with earlier work that characterizes parkour as curious play [2], highlighting that an interesting environment should allow for different types of uses. High variety enables experiencing high novelty, which is one of the core appraisals contributing to the emotion of curious interest [46].

3.2.2 Exploring and Finding Spots. The curious exploration view of parkour is further highlighted in that "exploring the city" was the most frequently indicated way of finding parkour spots (6 out of 9). Nevertheless, 4 participants – nearly half of our sample – indicated

Figure 3: The spot map maintained by the Helsinki parkour community as a Google My Maps page (<https://urly.fi/1hMx>). Spots are marked with pins, larger areas as polygons, and longer runs as red paths (e.g., an esplanade lined up with rocks to run on top of).

also using existing spot maps. This is not surprising given that Helsinki has a community-maintained online spot map with a good coverage of the best spots, as shown in Figure 3. The high usage of the Helsinki spot map indicates that our work should be useful for cities and regions that do not already have such maps. Other ways of discovering parkour spots included friends (4 participants) and social media (3 participants).

In addition to variety, participants highlighted the importance of appropriate challenges, related to both technique and creativity, e.g., *"A big variety of difficulty levels and possibility for progressions. Something for beginners and warm ups, low and easy but also something more challenging."*, *"Challenges that one can't beat right away, which motivates revisiting a spot once one's skills develop"*, *"Creative action and problem solving, self-expression. Simplicity that forces one to think and invent new ways to move."* Aesthetics and privacy were also mentioned: *"I prefer secluded spots. I like concrete but greenery makes a spot more enjoyable."*

3.2.3 Reasons for Enjoying Parkour. Common reasons cited for liking parkour in our sample are in line with earlier work:

- Interesting challenges (8 participants) *"I like the challenges that come with the sport (both physical and mental)"*
- Creativity (5 participants). *"I like being creative and being able to explore places by moving.", "Funny challenges, moving your body in a way you've never done before."*
- Freedom and autonomy (4) *"I like the sense of freedom I get when I train", "Competing only against yourself, no wrong or right ways to do stuff", "Parkour is easy to get back into after a break, because it's easy to adapt the difficulty level."*
- Community, fellowship (3 participants). *"I was originally drawn to the playfulness and social aspects of parkour.", "Fun people who make training fun, love to goof around, try stupid and silly things but still also train hard.", "I love how training is a combination of doing your best and pushing yourself as far as you can but also laughing more than anywhere else."*

The answers above highlight the importance of variety in desirable parkour spot qualities. A spot with highly varied geometry provides more opportunities for creative exploration and more freedom to select challenges appropriate for one's skill level. Large

Figure 4: An example of a playground specifically designed for parkour. Image courtesy of Lappset Group Ltd. Photographer: Antti Kurola.

spots with multiple shapes also support social training, as traceurs can practice multiple things in parallel.

3.3 Discussion

In summary, the results indicate that one should think of good parkour spots as (small) *playgrounds* consisting of versatile parkourable shapes, as opposed to just one or a few specific objects. This view is also supported by the recent emergence of playgrounds dedicated to or heavily inspired by parkour practice, e.g., Figure 4. The implication for the rest of this paper is that it is not enough that an automatic parkour spot detection system recognizes suitable geometry and objects such as stairs and rails; such a system must also understand how the interplay of the objects provides interesting and varied challenges.

4 SYSTEM

This section details our data collection and annotation approach, the machine learning architecture, and the visualizations developed for testing and evaluating the technology.

4.1 Data Collection

For training the network and validating system parameters, we used images loaded through the Google Street View API. We also considered using Flickr images tagged with the word "parkour", but these usually include parkour athletes in addition to spot geometry, which could confuse the learning. Two different sampling methods were used: 1) sampling a given area in a grid-like fashion, and 2) an informed approach where we obtained images of the pin locations of the community-maintained online Helsinki parkour spot map in Figure 3. Images of each map pin location were downloaded from 10 different angles.

The total number of training and validation images was 9,061. This comprises grid-sampled sets of 5003 images from Helsinki, Finland, 3252 images from Paris, France, and the informed sample of 806 images.

Finally, we collected a test data set with 585 images from a region in Espoo, Finland, which was not included in the other image sets. To test how well the system generalizes, we captured the test images ourselves while walking around the region, instead of downloading through the Google Street View API. This also allowed us to take multiple photos of parkour spots to test how sensitive the system is to different capture angles and distances.

4.2 Data Annotation

We recruited 5 parkour hobbyists to annotate the data, in addition to one of the authors who had 3 years of parkour experience. The annotators were recruited via the local parkour Facebook group and received two movie tickets each as compensation. The annotation was done in the browser-based LabelBox environment [19], which allows annotators to conveniently work wherever and whenever they choose.

Initially, the authors themselves tried labeling all common objects such as railings, poles, stairs, and low walls. However, this turned out to be very monotonous, tiring, and prone to errors such as missing one object out of many similar ones. During the annotation process, it also became clear that the most interesting images and geometry did not conform to the predetermined object classes. The interview study echoed this; based on the study, good parkour spots are not composed of single objects but instead of several that have suitable distances, placements, etc. There are no clear rules for defining what makes an interesting training spot.

In the end, we settled on a method with just two classes: very/definitely interesting and somewhat/maybe interesting. This makes the method fast and also easy to apply in other contexts, such as skateboarding, and also allowed for faster annotation of images. The first approach, which one annotator tested with 1092 images, required 14 seconds per image, on average. With the final approach, the same annotator only used 5s per image. The times spent by the other annotators were 7s, 4s, 4s, 5s, 16s, indicating that the average annotator can process almost 1000 images per hour, although one annotator took considerably more time than the others. The total annotation time of all 9k images was 17 hours, i.e., approximately two work days.

The annotators were instructed as follows:

Please mark interesting parkour spots/areas with rectangles encompassing all the geometry you would like to use when practicing (e.g., rails, stairs, walls, corners). For example, if there is a pole close to a rail for jump & swing practice, draw the rectangle such that it includes both the pole and the rail. If the pole is very tall, you don't need to include all of it, only the part that you would use in practice. There's no point in marking every common object like rails and poles. Only mark things that you would pay special attention to and/or be curious about when walking around the city. Please only consider the images, disregarding your prior knowledge you may have about any spots included in the images. For example, if you know there is a good spot but it is occluded in the image, don't mark anything.

Definitely interesting area: Something you would definitely want an automatic parkour map generator to notice and highlight.

Maybe interesting area: Something that you are less enthusiastic about or where it's unclear whether the spot is really useful. You'd like these areas to be highlighted on a parkour map, but omitting them would not be a big problem either.

Location	Images	Interesting	Interesting%
Helsinki, grid-sampled	5003	202	4.0%
Paris, grid-sampled	3252	217	6.7%
Helsinki, informed	806	359	44.5%
Espoo (test data)	585	132	22.6%

Table 1: Image annotation statistics. Interesting images are those with at least one maybe interesting or definitely interesting annotated area.

Table 1 summarizes the annotation statistics. The statistics indicate the randomly sampled Google Street View images only rarely contain interesting areas, and if one wants more positive training examples, it is useful to curate the training data, e.g., by utilizing existing spot maps.

We also considered other annotation schemes, such as labeling parkour spots based on their perceived challenge level or which parkour moves they would allow for, but ultimately decided against them due to our limited annotation capacity and because these qualities would likely be hard to reliably determine based on just a single 2D image.

4.3 Network Architecture and Training

To reduce the number of needed training images, we utilize transfer learning, i.e., a training curriculum where an initial learning task makes the final task easier [6]. Specifically, we started with a network pre-trained with generic image data and common object classes, and only finetuned the network with our own data.

We selected Facebook AI Research’s Detectron code repository [10] as the basis of our implementation, as it showed good results and supported our use case with minimal modifications to the original codebase: It already supported “freezing” the convolutional layers, a feature crucial to using transfer learning, and only a small edit of 17 lines to the code was required to allow the use of images with no annotations as negative training examples. This change was necessary to reduce the number of false positives from images that resemble a potential parkour spot but are not suitable due to the context the spot appears in, such as railings on highway bridges. Apart from this change, we only needed to modify the configuration files to adjust for our use case by, for example, defining the use of a pre-trained model and our custom dataset, and modifying the number of object classes being detected.

Of the pre-trained models included in Detectron, we selected the Faster R-CNN neural network architecture [38] with the ResNeXt-101-32x8d backbone, as this combination had fast training and inference times, as well as one of the highest box average precisions among all the models on the original task. The model was originally trained using images from the ImageNet-1k data set [47], meaning that it was primed to detect common object classes like people, cars, and bicycles from a large variety of images including urban imagery. Due to the relative similarity of the domains, we were fairly confident that this would be a good fit for our purposes.

In fine-tuning the network, we kept the convolutional feature extraction layers intact and only retrained the fully connected final layers. In other words, the network does not need to learn to extract features; it only needs to learn which feature combinations

represents the objects or geometries of interest. This reduces both the number of network weights to learn and the amount of training data needed. The Detectron implementation [10] supports “freezing” of convolutional layers out-of-the-box, and one only needs to edit a configuration file to change the number of object classes.

We trained the network for 180,000 minibatch iterations using a Tesla P100 GPU with 8 gigabytes of memory. One training run lasted for approximately 8 hours.

4.3.1 Early Stopping, Validation Metric. To prevent the network from overfitting to the limited data we used early stopping [11], i.e., we saved snapshots of the network during training and finally used the network that yielded the highest validation accuracy.

The Faster R-CNN network’s training objective function is based on how well the detected areas overlap with ground truth areas. However, as our goal is to detect suitable locations for parkour, the exact bounding boxes and overlaps are not significant. Thus, we computed the validation accuracy simply as the percentage of images where either 1) both the network’s output and the validation data had one or more interesting areas or 2) neither the network’s output or the validation data had interesting areas. Note that Faster R-CNN outputs a score in the range 0..1 for each detected area; in computing the accuracy, we only considered detections above a threshold value. We tested different threshold values between 0.59 and 0.99 with steps of 0.1.

4.3.2 Data Need. Modern deep learning achieves remarkable results in many computer vision tasks but often requires large amounts of data. Transfer learning should be able to reduce data requirements. We tested this by training the network with different amounts of training data ranging from 1000 to 7000 images. We had 9061 total images; we used a subset of 2061 images as the validation dataset for computing detection accuracy and determining the early stopping position.

4.3.3 Training Example Weighting. Our data annotations of definitely and maybe interesting areas were combined into a single-class detection problem with definitely interesting training examples having a higher weight. This had the benefit of focusing training efforts to detecting the most/definitely interesting areas and additionally helped to adjust for the imbalance between interesting and not interesting images in our dataset.

4.4 Data Visualization

To test and scrutinize the usefulness of the neural network outputs, we created both static and interactive visualizations.

Figure 5 shows examples of correctly and incorrectly detected areas from our test dataset. The top row of Figure 5 shows the most salient images, i.e., images sorted by the detection scores output by the network. Multiple images from the same spot have been pruned based on image GPS coordinates. The bottom row of Figure 5 shows randomly selected error images, i.e., images where either a human or the network detected something, but not both.

We also created prototypes of interactive visualizations, shown in Figure 1, Figure 6 and the supplemental video. In creating the visualizations, we first considered the visualizations that traceurs already create and use. The Helsinki parkour community has created a fairly detailed digital spot map as a custom Google My Maps

Figure 5: Example detection results with test data from a map region not used for training or hyperparameter tuning. Ground truth human annotations shown in green, neural network predictions in magenta. Top: Highest scoring detections featuring typical parkourable geometry such as low walls, rocks, and railings. Multiple images from the same spot have been pruned based on image GPS coordinates. Bottom: randomly selected error images, including 7 false negatives and one false positive.

page, shown in Figure 3. Although the map is already useful and nearly half of our interview respondents are using it, it has the following limitations:

- Most of the pin descriptions do not have images and also do not provide an easy way to show images, e.g., through linking to Google Street View.
- The basic view does not show any images to highlight particularly interesting spots or inspire one’s mind about parkour opportunities.
- The pin-based view would easily become too crowded if one would add hundreds of spots automatically identified from images.

Most online maps that we have found have been created with Google My Maps or similar platforms that suffer from the same limitations. As an example of how to overcome these limitations, Figure 1, Figure 6, and the supplemental video² show prototypes where the user can control the amount of spots shown by adjusting the detection score threshold with a slider. A thumbnail of the image closest to the pointer is shown, which allows quick browsing of the data. Clicking with the left mouse button enlarges the thumbnail for closer inspection. The prototype in Figure 6 also augments the map with a grid view of all thumbnails for each visible pin, sorted in decreasing order of spot probability.

5 EVALUATION

We evaluate the developed system and visualizations from three points of view:

- Quantitative evaluation: How accurate is the network in detecting interesting parkour spots?
- Qualitative evaluation: Looking at example results, what kind of spots does the network detect well, and what kinds of errors does it make?
- Feedback from the parkour community: Do parkour hobbyists see value in the results? What benefits and challenges do they identify?

5.1 Quantitative evaluation

Figure 7 shows the validation accuracy for each tested training data set size, averaged over three independent training runs with

Figure 6: A browser-based interactive spot map generated from our test dataset.

randomly initialized networks and training data subsets. The accuracy grows with the training dataset size, but there is also some random fluctuation. This is at least in part due to our small validation dataset of 2061 images, which increases the variance of the validation accuracy. The best validation accuracy over a single training run was 92.24%, which was reached after 150,000 training iterations, using a definitely interesting area weight 4, detection threshold 0.59, and 7000 training images.

A 100% accuracy is practically unreachable in our case, as the classification task is subjective and the annotators have differing opinions of what is interesting. The accuracy range in Figure 7 is also compressed because our training data: Due to the rarity of interesting areas in the data, the naïve approach of always predicting that there are no interesting areas yields an accuracy of 91.3%. Correcting for the class imbalance is not motivated in our case, as the imbalance is not due to incorrect training data collection. Instead, the imbalance reflects the real-life sparsity of interesting areas, and a good parkour spot detector should indeed reject most images and only highlight the rare interesting ones.

²<https://youtu.be/vFCcXTicqNE>

Overall, we acknowledge that the accuracy leaves room for improvement, but as elaborated below, visualizing and exploring parkour spots is a low-risk activity where some errors can be tolerated.

5.2 Qualitative evaluation

As shown in Figure 5, the highest-scoring detections are sensible and the network correctly identifies clearly parkourable geometry such as rails, walls, and stairs. However, the errors made by the system are perhaps more interesting in analyzing what it understands and what should be improved in future systems.

5.2.1 What kinds of errors does the system make? The false positive image at the top-right corner Figure 5 illustrates the subjective and imprecise nature of the human annotations. The same spot also shows up as a false negative (bottom row, 4th from the left), but that image is actually a data annotation error; the annotator mistook a shadow for a platform. The spot is also confusing because it features a large raw concrete shape, which usually draws a traceur's eye; concrete shapes are stable, durable, and provide excellent grip, which makes them ideal for parkour.

Most of the errors in Figure 5 are false negatives, i.e., where a human marked an area but the network did not detect anything. The amount of false negatives can be decreased by using a lower detection threshold, as shown in the supplemental video. On the other hand, this also increases the amount of false positives, which may cause clutter in data visualization. Ultimately, the optimal detection threshold depends on the visualization or application.

The 3rd error image from the left in Figure 5 is a spot that is correctly detected in other images, but probably missed in this case because of small scale. The 5th image shows a pyramid shape that is unique to the test dataset. Although slanted walls are interesting for sliding and climbing, they are very rare in the training data. Sometimes, it also seems that the network understands geometry

but not material affordances; the false positive on bottom-right corner of Figure 5 would be an interesting spot if the windowsill was actually suitable for standing and the round shrubs were rocks instead. It is possible that the material of the shrubs is unclear because of the heavy shadows.

5.2.2 Data Annotation Quality. The errors further highlight the subjective nature of the training data. Only images 1, 3, 5, and 7 (from the left) of Figure 5 are clear false negatives. Image 2 is less clear, as the vegetation surrounding the rocks somewhat impairs both visual detection and parkourability. As discussed, image 4 is an annotation error, and image 6 is likewise an error or a borderline case. While the low wall can be used for vault practice, the single shape does not provide much variety, and the data annotator remembers questioning whether he should annotate it or not.

We also spotted at least one training image where an annotated area was so small that the annotator probably marked it based on their experience of the area instead of what is actually visible in the image. We advised the annotators against this, but mistakes are made easily in the monotonous annotation process.

5.3 Feedback from the Parkour Community

We solicited feedback from the local parkour community in two ways. First, we posted spot detection results to the Parkour Helsinki Facebook group for commenting. Second, we recruited traceurs to evaluate the browser-based spot map prototype we generated from our data, shown in Figure 6 and on the supplemental video at 01:52.

5.3.1 Feedback on results posted on social media. We generated a sequence of 10 top predictions and 10 error images using the validation dataset and posted them to the Parkour Helsinki Facebook group, asking for feedback. In total, seven traceurs commented on the images. The respondents were not compensated. The comments pointed out that traceurs can have highly differing opinions about what is considered an interesting area. Two respondents commented that out of the 10 false negatives, only 4 were real errors and the rest of the images did not really contain anything interesting. As a caveat, it was pointed out that similar to one's own parkour gaze, visual inference is limited and one cannot be certain of a spot's qualities without actually trying it out. For example, it may turn out that the distances between objects are too short or long, or the surface materials are too slippery.

Overall, feedback was positive, and the traceurs identified many of their common training spots among the top detections. The responses also indicated that the system gives useful suggestions of spots to check out. For future work, it was requested that the system would also work with aerial images, e.g., to recognize rock formations and walls or ruins.

5.3.2 Evaluation of the browser-based interactive map. As the social media feedback consisted of mostly brief positive comments with little information, we recruited four traceurs for more in-depth testing and insights. Due to the COVID-19 lockdown which prevented meeting with user study participants, we had participants use the browser-based prototype in Figure 6 in their homes, while sharing their screens with a researcher over a Zoom video call. For comparison, participants also used the community-created Google Map in Figure 3. Ages of the practitioners ranged from 26 to 40 years old.

Figure 7: Validation accuracy as a function of the number of training images, averaged over three independently trained networks. The red line depicts null accuracy, i.e., the accuracy achieved by always predicting no interesting areas. The null accuracy is high due to the rarity of interesting parkour spots in random images. Additional training images increase the accuracy.

Overall parkour experience of the practitioners ranged from 8 to 20 years. The order of map interfaces compared was counterbalanced. The study duration was approximately 30 minutes. Participants were provided 10 EUR for their participation.

For each tested interface, participants were instructed to:

"Imagine that you'd like to practice some parkour, but you are bored with your usual spots. You'd like to find some new spots to train at and explore what some city area can offer. Please try using the interface to accomplish that, narrating your thoughts and feelings as you go. You should use at most 5 minutes for this."

Participants could ask questions during use of each tool. Experimenters also aided in finding certain aspects of the tools, such as routes on Google My Maps and the adjustable spot detection threshold for our neural network.

Once participants felt confident with using the tool, they were asked:

- Do you think this tool can help you find interesting training spots? Why?
- What is good about this tool?
- What should be improved about this tool?

After answering these questions, the participant tried out the second tool and answered those questions again. After the participant was familiar with the usage of both tools, they were asked:

- How would you compare the two tools you tried?
- Which tool would you prefer and why?

Participants highlighted the simplicity of our system as both a pro and con. On the one hand, it was easier for participants to find locations, and unlike the Google My Maps tool, they were not provided with information overload. It was easier for participants to pick up and use right away, whereas the Google My Maps tool had a steeper learning curve, given its various features. On the other hand, the simplicity of the tool meant that several features participants enjoyed in the Google My Maps tool were not available in our tool. For instance, our tool lacks descriptions of the areas and the ability to zoom in on images. This meant that participants were not entirely sure why some spots were chosen by the system.

To improve the Google My Maps tool, participants suggested a filtering system by difficulty, as some spots may be more suited to expert-level traceurs, while others may be better suited for beginners. Additional filtering was suggested to filter by legality, as some spots on the Google map required jumping across rooftops, which is questionable on its legality.

To improve our tool, participants suggested a rating and feedback system from the community and the ability to overlay terrain imagery. In addition, participants requested knowledge of what parameters the system uses to select spots, as well as the ability to adjust these parameters.

One of the main problems of the Google My Maps interface we addressed in our own design was the lack of visual information. Indeed, the Google My Maps tool was negatively received by participants regarding the lack of pictures of the area. In the Google My Maps tool, community members can add pictures or even video of each training location, but it is much easier to add a simple description. Thus, most training spots on the Google My Maps tool

consist primarily of descriptions with very few pictures and even fewer videos.

6 DISCUSSION

Overall, our results indicate that it is possible to combine computer vision and street-level imagery to automatically detect urban physical activity opportunities such as parkour spots. Furthermore, using transfer learning makes this feasible with only a modest amount of data collection and annotation. We have also demonstrated the use of the generated data in building spot exploration and visualization tools, in order to support the curious exploration view of parkour highlighted by both previous research and our interview study, and to overcome the lack of visual information we have identified in existing spot maps.

Beyond the already mentioned possibility for a parkour-themed variant of Pokémon Go [32], it is easy to envision other applications for our data, e.g., automatically generating jogging paths that also visit parkour spots where one has to complete one or more exercises. A recent example in this vein is provided by Cityspotting [36], a gamified urban exploration app that motivates players to visit new locations and perform exercises such as hopping on seaside rocks. The app was in development during our spot map evaluation and the founder happened to be one of the participants who tested our prototype. He says that understanding urban geometry and its exercise affordances is crucial to their product, and our data gave him valuable insights into parts of the city that he was not familiar with.

As our participants asked for a way for the parkour community to provide feedback and ratings, a hybrid community-created and machine learning-based system might provide the best of both worlds. This would also align with the need for citizen-centered and inclusive processes of city-making [44], and the vision of future hybrid cities with immersive and collaborative digital layers or "mirror worlds", where citizens can become content creators [53]. User-generated content would also mitigate an inherent limitation we encountered in Google Street View imagery: Some parkour spots are not visible from the street. Figure 8 shows examples of this.

Ultimately, one might envision an open source and open data ecosystem that extends present tools such as OpenStreetMap through combining interaction design and machine learning. People interested in urban play and exercise could have efficient mobile interfaces and tools, not just for discovering locations and content, but for building a shared playground and community through contributing images, video, 3D scans of the environment, and associated data such as spot annotations or game levels (e.g., parkour flows or street workout challenges). Machine learning and computer vision could then automatically filter and anonymize the data (e.g., blurring faces like in Google Street View) and generate new content. Examples of the latter include generalizing human annotations to new data like in this paper, providing recommendations ("people who liked this spot/challenge also liked..."), or rendering AR content such as a virtual character or recorded "ghost" of another player providing a follow-the-leader challenge. Through combining mobile phone movement sensors, computer vision, and dedicated fitness tracking devices, one could also automatically collect leaderboard metrics

Figure 8: Two examples of how a street view may provide limited or no visibility of good parkour spots. Both illustrated spots are frequented by Helsinki traceurs, allowing varied practice of balancing, precision jumping, running up or vaulting over the low walls, falling and rolling on the sand etc.

such as the total time balanced on rails or total wall run steps per week and per spot, either for an individual user or collaboratively for a neighborhood.

As a downside, technological augmentation might restrict one's creativity and hinder the development of one's non-augmented parkour vision. It may be that technology tools for parkour spot discovery pose the danger of further severing our ability to be present and in immediate connection with our surroundings, as opposed to interacting through devices and apps. On the other hand, traceurs already use and create digital spot maps, and not every hobbyist has the time to explore and discover everything themselves. The effect of technological augmentation also depends on the specifics of the implementation. For example, Malinverdi et al. [24] found that a projected version of an Augmented Reality interface promoted more direct engagement with the environment than a screen-based version.

7 LIMITATIONS

Our data is imperfect in that we only employed a single annotator per image. Therefore, we cannot make conclusions of the biases and variance of the annotations. As elaborated in Section 5.2.2, the annotations are inevitably subjective and annotators may have also used their prior knowledge in addition to purely visual inspection when making decisions, which could make it harder to train a computer vision system on the data. Collecting a more extensive and higher quality dataset with multiple annotators for each image remains as future work.

Street level imagery might not be readily available or up to date for some regions. For example, there was an area in the Helsinki region that contained multiple parkour spots marked on the community map, but the area is poorly covered by Google Street View

images, and there were also some areas where spots were marked but images were almost a decade old. The latter can especially be a problem when automatically detecting potential spots, as new construction work might render old spots unusable and create new ones that cannot be discovered before the images are updated.

Further limitations of our work are that complete beginners might not understand the affordances of the spots recognized by our system, and we have not extensively prototyped and tested what kind of visualizations and Geographic Information Systems could be created using the neural network data. In future work, we aim to investigate both interactive parkour spot maps and see-through Augmented Reality "parkour vision" that would not only highlight interesting areas, but also show animations of movements to practice. The latter requires in-depth knowledge of the 3D geometry; fortunately, this is increasingly available through large scale photorealistic 3D "digital twin" city models such as Virtual Helsinki [23].

8 CONCLUSION

We have presented a two-part study: 1) An interview study about what kind of parkour spots are interesting and how parkour hobbyists find them, and 2) an experiment in training a deep neural network for the novel problem of automatically detecting interesting parkour spots from urban images which are easily available in large quantities through services like Google Street View and Flickr. The results were evaluated both quantitatively and qualitatively.

Our work provides a new tool for discovering and understanding physical activity opportunities in one's everyday environment, which should be valuable for researchers and practitioners of fields like exergame and urban design. The feedback from our local parkour community has been very positive, and we look forward to using our system to generate online spot maps of new cities, in order to help people find meaningful exercise environments and challenges.

More generally, our work illustrates how open source machine learning tools are maturing rapidly and can be easily applied, enabling new research opportunities. We demonstrate that using a transfer learning approach—i.e., fine-tuning an off-the-shelf pre-trained network with a small custom dataset—can keep the amount of data collection and annotation feasible, even if the annotations require expert knowledge and are thus challenging to crowdsource in vast amounts. We also show how testing and iterating on the data annotation scheme can enable considerable savings in annotation work. On the other hand, our work highlights the limitations of street view imagery, demonstrating hidden playgrounds about which a street view only provides subtle hints. In future work, this could be addressed by augmenting street view imagery with custom crowdsourced image data.

ACKNOWLEDGMENTS

This work utilized computational resources provided by the Aalto Science-IT project. We are also deeply grateful to the Helsinki parkour community for welcoming this project and participating in the interviews, data annotation, and the evaluation of the results.

REFERENCES

- [1] Lieven Ameen and Sirpa Tani. 2012. Everyday aesthetics in action: Parkour eyes and the beauty of concrete walls. *Emotion, Space and Society* 5, 3 (2012), 164–173.
- [2] Lieven Ameen and Sirpa Tani. 2012. Parkour: Creating loose spaces? *Geografiska Annaler: Series B, Human Geography* 94, 1 (2012), 17–30. <https://doi.org/10.1111/j.1468-0467.2012.00393.x>
- [3] Josh Andres, Julian de Hoog, and Florian 'Floyd' Mueller. 2018. I had super-powers when eBike riding Towards Understanding the Design of Integrated Exertion. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play*. ACM, 19–31.
- [4] Adrian E Bauman, Rodrigo S Reis, James F Sallis, Jonathan C Wells, Ruth JF Loos, Brian W Martin, Lancet Physical Activity Series Working Group, et al. 2012. Correlates of physical activity: why are some people physically active and others not? *The lancet* 380, 9838 (2012), 258–271.
- [5] Nathaniel Bavinton. 2007. From obstacle to opportunity: Parkour, leisure, and the reinterpretation of constraints. *Annals of leisure research* 10, 3-4 (2007), 391–412. <https://doi.org/10.1080/11745398.2007.9686773>
- [6] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*. ACM, 41–48.
- [7] Nadia Bianchi-Berthouze. 2013. Understanding the role of body movement in player engagement. *Human-Computer Interaction* 28, 1 (2013), 40–75.
- [8] Michael L Booth, Adrian Bauman, Neville Owen, and Christopher J Gore. 1997. Physical activity preferences, preferred sources of assistance, and perceived barriers to increased activity among physically inactive Australians. *Preventive medicine* 26, 1 (1997), 131–137.
- [9] Paul A Estabrooks, Rebecca E Lee, and Nancy C Gyurcsik. 2003. Resources for physical activity participation: does availability and accessibility differ by neighborhood socioeconomic status? *Annals of behavioral medicine* 25, 2 (2003), 100–104.
- [10] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. 2018. Detectron. <https://github.com/facebookresearch/detectron>.
- [11] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [12] Perttu Hämäläinen, Joe Marshall, Raine Kajastila, Richard Byrne, and Florian Floyd Mueller. 2015. Utilizing gravity in movement-based games and play. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*. ACM, 67–77.
- [13] Kotaro Hara, Vicki Le, and Jon Froehlich. 2013. Combining crowdsourcing and google street view to identify street-level accessibility problems. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 631–640.
- [14] Kotaro Hara, Jin Sun, Robert Moore, David Jacobs, and Jon Froehlich. 2014. Tohme: detecting curb ramps in google street view using crowdsourcing, computer vision, and machine learning. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 189–204.
- [15] Katherine Isbister and Florian "Floyd" Mueller. 2015. Guidelines for the design of movement-based games and their relevance to HCI. *Human-Computer Interaction* 30, 3-4 (2015), 366–399.
- [16] Mads Møller Jensen, Majken Kirkegaard Rasmussen, and Kaj Grønbaek. 2014. Design sensitivities for interactive sport-training games. In *Proceedings of the 2014 conference on Designing interactive systems*. ACM, 685–694.
- [17] Raine Kajastila, Leo Holsti, and Perttu Hämäläinen. 2016. The augmented climbing wall: high-exertion proximity interaction on a wall-sized interactive surface. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. ACM, 758–769.
- [18] Myron W Krueger, Thomas Gionfriddo, and Katrin Hinrichsen. 1985. VIDEO-PLACE—an artificial reality. In *ACM SIGCHI Bulletin*, Vol. 16. ACM, 35–40.
- [19] LabelBox. [n.d.]. LabelBox image annotation system. Online service, <https://labelbox.com>.
- [20] Christian Lander, Frederik Wiehr, Nico Herbig, Antonio Krüger, and Markus Löchtfeld. 2017. Inferring landmarks for pedestrian navigation from mobile eye-tracking data and Google Street View. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2721–2729.
- [21] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
- [22] Massimo Leone. 2012. The semiotics of parkour. *Ars Semeiotica* 35, 1-2 (2012).
- [23] Zoan Ltd. 2019. Virtual Helsinki. <https://www.virtualhelsinki.fi/>, Accessed: 14 Sep 2019.
- [24] Laura Malinverni, Julian Maya, Marie-Monique Schaper, and Narcis Pares. 2017. The World-as-Support: Embodied Exploration, Understanding and Meaning-Making of the Augmented World. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). ACM, New York, NY, USA, 5132–5144. <https://doi.org/10.1145/3025453.3025955>
- [25] Elena Márquez Segura, Annika Waern, Jin Moen, and Carolina Johansson. 2013. The Design Space of Body Games: Technological, Physical, and Social Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). ACM, New York, NY, USA, 3365–3374. <https://doi.org/10.1145/2470654.2466461>
- [26] Joe Marshall, Frank Loesche, Conor Linehan, Daniel Johnson, and Bruno Martelli. 2015. Grand Push Auto: A car based exertion game. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*. ACM, 631–636.
- [27] Rich McEachran. 2017. These 7 playful ideas can bring fun to the city. <https://www.smart-magazine.com/urban-gamification-solutions/>.
- [28] Florian Mueller, Stefan Agamanolis, and Rosalind Picard. 2003. Exertion interfaces: sports over a distance for social bonding and fun. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 561–568.
- [29] Florian 'Floyd' Mueller, Darren Edge, Frank Vetere, Martin R Gibbs, Stefan Agamanolis, Bert Bongers, and Jennifer G Sheridan. 2011. Designing sports: a framework for exertion games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2651–2660.
- [30] Florian 'Floyd' Mueller, Richard Byrne, Josh Andres, and Rakesh Patibanda. 2018. Experiencing the Body As Play. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). ACM, New York, NY, USA, Article 210, 13 pages. <https://doi.org/10.1145/3173574.3173784>
- [31] United Nations. [n.d.]. Sustainable Development Goals. <https://www.un.org/sustainabledevelopment/sustainable-development-goals/>, accessed 3 Sep 2019.
- [32] Niantic. 2016. *Pokémon Go*. Game [iOS, Android].
- [33] Anton Nijholt. 2017. Towards playful and playable cities. In *Playable cities*. Springer, 1–20.
- [34] Anton Nijholt. 2020. Playful Introduction on "Making Smart Cities More Playable". In *Making Smart Cities More Playable*. Springer, 1–22.
- [35] Alice O'Grady. 2012. Tracing the city—parkour training, play and the practice of collaborative learning. *Theatre, dance and performance training* 3, 2 (2012), 145–162. <https://doi.org/10.1080/19443927.2012.686450>
- [36] Fennoscandic Simulation Systems Oy. [n.d.]. Cityspotting. Online service, <https://cityspotting.fi>.
- [37] Christopher Rawlinson and Mirko Guaralda. 2011. Play in the city: Parkour and architecture. In *The First International Postgraduate Conference on Engineering, Designing and Developing the Built Environment for Sustainable Wellbeing*. Queensland University of Technology.
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [39] Nina Runge, Pavel Samsonov, Donald Degraen, and Johannes Schöning. 2016. No more autobahn!: Scenic route generation using googles street view. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 147–151.
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115, 3 (2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [41] Johnny Saldaña. 2015. *The coding manual for qualitative researchers*. Sage.
- [42] James F Sallis, Melbourne F Hovell, C Richard Hofstetter, John P Elder, Mimi Hackley, Carl J Caspersen, and Kenneth E Powell. 1990. Distance between homes and exercise facilities related to frequency of exercise among San Diego residents. *Public health reports* 105, 2 (1990), 179.
- [43] Stephen John Saville. 2008. Playing with fear: Parkour and the mobility of emotion. *Social & cultural geography* 9, 8 (2008), 891–914. <https://doi.org/10.1080/14649360802441440>
- [44] Ben Schouten, Gabriele Ferri, Michiel de Lange, and Karel Millenaar. 2017. Games as strong concepts for city-making. In *Playable Cities*. Springer, 23–45.
- [45] Andres Sevtuk. 2013. Networks of the built environment. *DeJ coding the City—How Big Data Can Change Urbanism* (2013), 144–59.
- [46] Paul J Silvia. 2008. Interest—The curious emotion. *Current Directions in Psychological Science* 17, 1 (2008), 57–60.
- [47] Princeton University Stanford Vision Lab, Stanford University. [n.d.]. ImageNet statistics. <http://image-net.org/about-stats>, accessed 19 Sep 2019.
- [48] Jeffrey Star and John E Estes. 1990. *Geographic information systems: an introduction*. Vol. 303. Prentice Hall Englewood Cliffs, NJ.
- [49] Marlene K Tappe, Joan L Duda, and Patricia M Ehrnwald. 1989. Perceived barriers to exercise among adolescents. *Journal of School Health* 59, 4 (1989), 153–155.
- [50] Mattia Thibault. 2019. Towards a typology of urban gamification. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- [51] Annika Waern, Elena Balan, and Kim Nevelsteen. 2012. Athletes and Street Acrobats: Designing for Play As a Community Value in Parkour. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). ACM, New York, NY, USA, 869–878. <https://doi.org/10.1145/2207676.2208528>
- [52] Xiao Xiao and Hiroshi Ishii. 2016. Inspect, Embody, Invent: A Design Framework for Music Learning and Beyond. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). ACM, New York, NY, USA, 5397–5408. <https://doi.org/10.1145/2858036.2858577>
- [53] Johanna Ylipulli, Jenny Kangasvu, Toni Alatalo, and Timo Ojala. 2016. Chasing Digital Shadows: Exploring future hybrid cities through anthropological design fiction. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. 1–10.