Constellation: a Multi-User Interface for Remote Drone Tours

Jiannan Li
jiannanli@dgp.toronto.edu
Department of Computer Science, University of Toronto
Toronto, Ontario, Canada

Maurício Sousa
mauricio@dgp.toronto.edu
Department of Computer Science, University of Toronto
Toronto, Ontario, Canada

Ravin Balakrishnan
ravin@dgp.toronto.edu
Department of Computer Science, University of Toronto
Toronto, Ontario, Canada

Tovi Grossman
tovi@dgp.toronto.edu
Department of Computer Science, University of Toronto
Toronto, Ontario, Canada

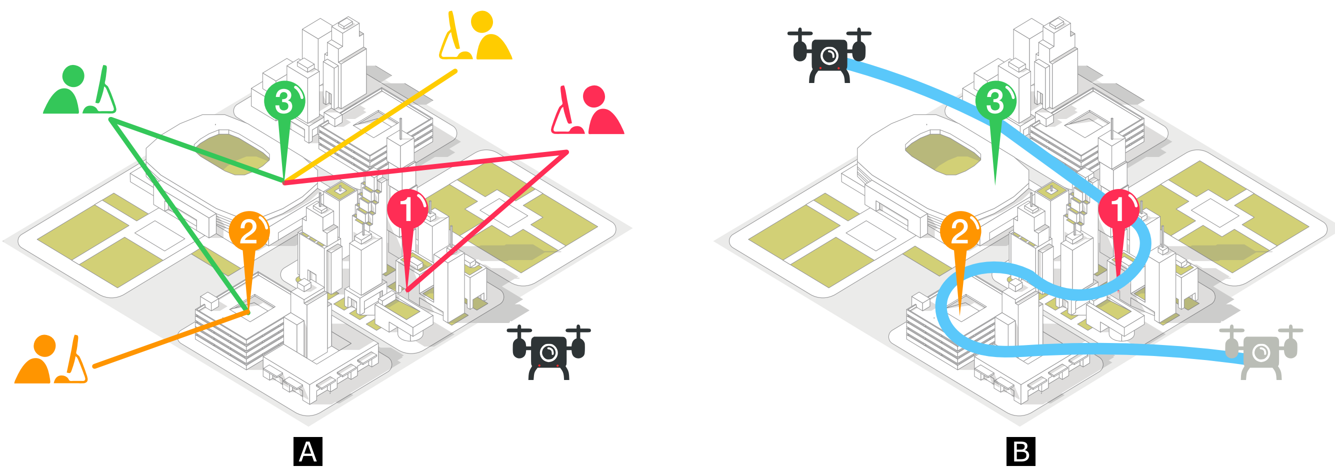

Figure 1: Multiple users on a remote drone tour using the Constellation system share the control of a camera drone. (A) Each user can specify points of interest and upvote other users' point of interest. The three points of interests have 1, 2, 3 votes, respectively, from 4 users. (B) Constellation plans a drone path that minimizes the total extra waiting time of all users.

ABSTRACT

Remotely controlled camera drones can support live, dynamic, and interactive virtual tours for travelers to overcome distance, expense, and health barriers. Yet, assigning one drone to one traveler may incur unnecessary waste of resources, and an abundance of concurrent drones raises safety concerns. While sharing the input and output of a single drone among multiple concurrent users can alleviate these limitations, standard control sharing protocols, such as turn-taking, are often inefficient. We present Constellation, a multi-user drone control system that synthesizes diverse user goals and generates efficient flight paths for the group. It supports point-of-interest specification on both static 3D environmental maps and live camera views. The generated paths minimize all users' total extra waiting time. A web-based study with 16 participants show that Constellation could help groups navigate to their points-of-interest faster in comparison to the turn-taking baseline.

CCS CONCEPTS

• **Human-centered computing** → **Interaction techniques**; **Interactive systems and tools**.

KEYWORDS

Navigation, Drone Interaction, Interaction Design

ACM Reference Format:
Jiannan Li, Maurício Sousa, Ravin Balakrishnan, and Tovi Grossman. 2021. Constellation: a Multi-User Interface for Remote Drone Tours. In *Proceedings of the 9th International Conference on Human-Agent Interaction (HAI '21), November 9–11, 2021, Virtual Event, Japan.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3472307.3484685

HAI '21, November 9–11, 2021, Virtual Event, Japan

ACM ISBN 978-1-4503-8620-3/21/11.
https://doi.org/10.1145/3472307.3484685

1 INTRODUCTION

People enjoy traveling to alleviate daily stress and broaden their world-views [11]. Although the concept of travel is primarily associated with going to remote places, from the 19th-century travel journals to modern Twitch[1] travel live streams [32], virtual tourism

[1]Twitch: https://www.twitch.tv

enables people to perceive traveling experiences with satisfying levels of authenticity, despite not being there [35]. While virtual tours often supplement later or earlier physical travel experiences [7], they can also be vital for those who cannot travel because of time, cost, and health reasons [24, 35]. For example, the environment and geography education community employ virtual field trips to offer accessible learning experiences to financially or physically disadvantaged students [41, 46, 47].

Users consuming virtual tour experiences often prefer live video streams over pre-recorded videos due to their higher level of perceived authenticity [32]. The robotics community has explored live video tours through remotely operated robots [37, 52], which offer the additional benefit of direct control over what to see. Camera drones offer unique mobility and perspective benefits [33, 34, 38, 50], as they are less constrained by barriers [2] and provide informative vantage points [38]. Recent advances in autonomous drone navigation allow untrained users to remotely pilot camera drones with ease [9, 17, 26, 48, 49, 53]. While most remote drone navigation interfaces follow a one-user-one-robot paradigm, having one visitor control one drone may result in a large crowd of drones at popular tourism sites. Alarms have been raised about the energy waste of an abundance of drones [14, 31] and the associated safety [3, 39] and privacy [54] risks for local communities. A mechanism for multiple users to share the control of a drone can potentially mitigate these concerns and make remote drone touring more friendly to the environment and local communities.

In this paper, we present Constellation, a novel multi-user interface for remote drone navigation. Constellation is based on a simulated drone in virtual environments for prototyping interactions and conducting preliminary evaluations. A significant challenge for multi-user drone control in remote tours is how to accommodate the diversity in users' goals while maintaining a reasonable level of time efficiency. While turn-taking is a standard method for handling heterogeneous user goals, previous work has noted its relative inefficiency compared to an interface that supports multiple synchronous inputs [5, 23]. Prior research in multi-user agent control [19, 27, 44] applied voting schemes from social choice theories [6] to combine multiple synchronous inputs into a single command. These approaches use majority consensus to produce a reliable or socially desirable aggregate input, but they have two limitations for control sharing in drone tours. First, they either discard or distort minority preferences. Second, they do not consider factors other than preferences, such as user goal locations, therefore producing less efficient plans.

Research in human-robot interaction has studied modeling user utility for robot control, but mainly in single user scenarios [10, 21]. We apply user utility optimization to devise a novel multi-user control sharing mechanism that accommodates diverse inputs and balances considerations for user preferences and for physical realities. More specifically, Constellation finds drone paths that minimizes the total extra waiting time of all users. Multiple people can use the Constellation interface to simultaneously explore a remote environment and specify their points of interest on 3D environment maps and camera live streams. Users can also efficiently express their preferences by upvoting or modifying other peoples' points of interest. The system constantly collects and synthesizes their preferences to generate drone flight plans that accommodate all the inputs (Figure 1). We also provides tour and environment information to help users stay oriented during the tour.

We compared Constellation with a turn-taking baseline in a group navigation task. The results from the two online study sessions, each with eight concurrent users, validated the efficiency benefits of our approach. We discuss future work on adapting Constellation for real-world drone tours, and on designing multi-user robot interfaces for other domains.

2 RELATED WORK

Our work builds on previous research on human-drone interfaces and multi-user agent control. Prior research has explored the versatile roles of drones [1, 8, 20, 25, 26]. To alleviate pilots' burden during manual control [53], recent drone interfaces incorporate designs to support stronger situation awareness from a second drone [49], object-centric high-level commands [13, 30], and direct-manipulation control through augmented reality [9]. Building on prior human-drone interface designs, we explore drones' new role as a shared resource in virtual tourism. While our current implementation runs only in virtual environments, its workflow and interface design builds on established real drone videography systems [17, 25, 43, 53] for practicality.

Researchers in internet robotics [19, 45] and online gaming [29, 42] have studied mechanisms for group control of a single agent for resource sharing. Meanwhile, the human computation community has experimented with various methods for aggregating commands to achieve robust crowd robot control [27, 44]. Their approaches typically leverage the consensus among user inputs to produce a socially desirable aggregate command or filter out noise. To do so, they apply various voting schemes from social choice theories [6], including plural [19, 42], weighted [27], and Borda [44] voting mechanisms. As an alternative perspective, Higuera et al. [22] presents a collective path planner that prioritizes shorter paths but does not respect consensus. We propose a control sharing mechanism that accommodates all inputs and achieves user goals efficiently by considering both group consensus and goal locations.

3 CONSTELLATION: A MULTI-USER DRONE INTERFACE

Constellation is a multi-user remote drone touring interface with utility-optimization as its underlying control sharing mechanism. It allows a user to join online with a group of peer travellers to experience a remote environment through a live drone camera feed. Users can control the drone together to explore the environment interactively, reacting to ongoing changes such as live performances encountered in remote sightseeing.

The user interface of Constellation has two main modes, the Camera Stream mode and the 3D Map View mode. Both modes incorporate a key widget named View Queue (Figure 2, next page). Users control the drone through sending *view requests*.

In *Camera Stream* mode, users see a full-screen live stream from the drone camera (Figure 2 left). Our prototype shows the virtual camera feed from the simulated drone camera in high fidelity. Scrolling the mouse wheel down on the camera stream turns on the *3D Map View* mode, which shows a 3D map of the surrounding environment (Figure 2 right). We segmented the meshes in the 3D

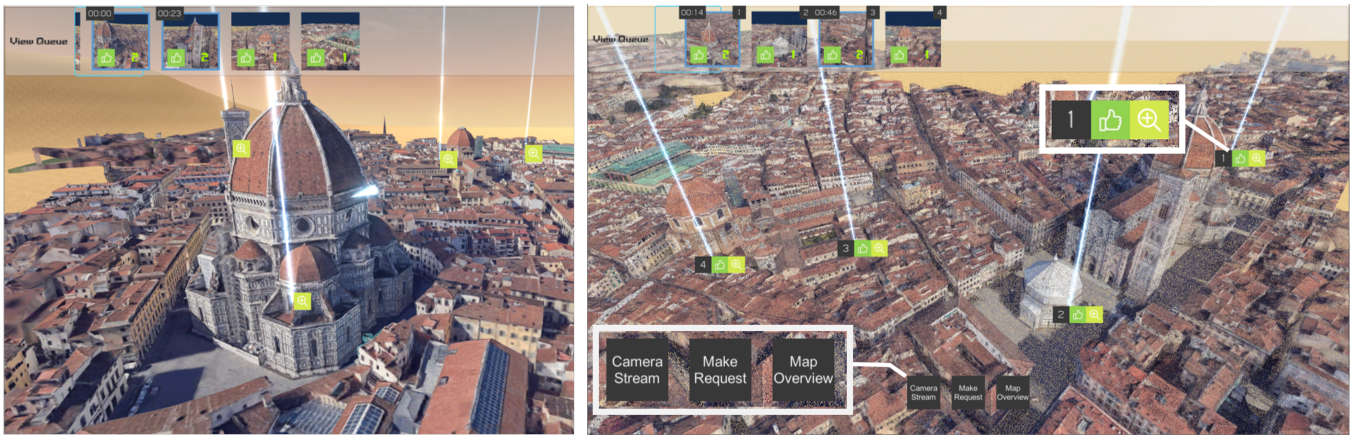

Figure 2: Overview of the Constellation interface. *Left*: the Camera Stream mode. *Right*: the 3D Map View mode with the upvote and preview buttons in the callout. The View Queue is at the top in both sub-figure.

map into individual objects. In this view, users can roam freely on the map using standard virtual camera manipulation techniques. The user can return to the camera stream mode by clicking a button (Figure 2, right). In our current implementation, the 3D Map View is processed with a point-cloud-effect shader to simulate real-world 3D scans. The *View Queue* sits at the top of the screen on both modes (Figure 2). It displays a queue of current view requests, arranged by their position in the future flight path.

4 PATH PLANNING FOR DIVERSE USER GOALS

After users specify their goals via the interface, Constellation generates a path using the utility optimization framework. The path is updated constantly as new requests come in. Now we introduce the utility optimization framework.

4.1 Path Planning as Utility Optimization

We can describe the formulation of collective drone path planning as a utility optimization problem with three key components, the decision variables and constraints, the user utility function, and the objective function.

4.1.1 Decision Variables and Constraints. Formally, the path planning problem concerns with arranging the order of N view requests. It can be viewed as a variation of the Travelling Salesman Problem (TSP) [4]. We model it as an assignment problem, where N view requests need to be assigned to N 'slots' that constitute the path. We chose this formulation instead of common TSP formulations as it enables a concise expression of our utility function.

To simplify the formulation that follows, we treat the starting position of the drone as a view request v_1. Given $N-1$ more view requests $v_2, \dots v_N$ and N slots on the path, the assignment can be expressed as the *planning vector* $\mathrm{x} = (x_{11}, x_{12}, \dots, x_{vv}) \in \mathbf{X}$, where

$$x_{ij} \in \{0,1\}\ \forall i \in \{1, \dots N\}, j \in \{1, \dots N\}, \text{ such that}$$

$$x_{ij} = \begin{cases} 1 & \text{if view request } v_i \text{ is assigned to slot } j \\ 0 & \text{otherwise} \end{cases}$$

The set $\mathbf{X}$ contains all assignment choices. A feasible path assignment is constrained by

$$\begin{gathered} \sum_{j=1}^{N} x_{ij} = 1,\ \forall i \in \{1, \dots N\} \quad \text{and} \\ \sum_{i=1}^{N} x_{ij} = 1,\ \forall j \in \{1, \dots N\} \quad \text{and} \quad x_{11} = 1 \end{gathered} \tag{1}$$

4.1.2 User Utility Function. As the first step in exploring user utility in shared drone control, our utility model has a simple structure. It is widely recognized by marketing research that people's queuing experience deteriorates as their time in the queue exceeds their expected waiting time [12, 40, 51]. We assume that a user's expected waiting time T for a view request is the time needed to visit this location if the drone would take the shortest path to fulfill all of her view requests. We then use the difference between the actual time t spent to fulfill a view request and the expected waiting time T as the cost in utility calculation. More specifically, when the drone fulfills a view request, any user who has created or voted for this request gains a fixed benefit of value g and pays the cost of their extra waiting time $t - T$. That is, the utility π of a user for a single view request being fulfilled can be written as $\pi = g - (t - T) = g - t + T$. T is a constant that depends on the drone position at the time of planning and the locations of the view requests.

The preferences of any user u_i in the user set U with respect to all N view requests V can be represented with a N dimensional vector $\mathrm{p_i} = (p_{i1}, p_{i2}, \dots, p_{iN})$, in which $p_{ij} = 1$ if user u_i has created or voted for view request v_j otherwise 0. We call a view request v_j a *relevant request* of user u_i if $p_{ij} = 1$. We further define a N by N distance matrix D whose element d_{ij} denotes the distance between two view requests v_i and v_j and assume a constant drone speed of 1. On a path where two consecutive slots q and $q+1$ are assigned to the view requests v_m and v_n, respectively, as user u_i follows the drone moving from slot q to slot $q+1$, her incurred cost is the time to travel between the two view requests d_{mn} multiplied by r_{iq}, the number of her unfulfilled view requests after visiting slot q. r_{iq} can be computed by

$$r_{iq} = \sum_{k=1}^{N} p_{ik}(1 - \sum_{j=1}^{q} x_{kj}),\ \forall i \in \{1, \dots N\},\ \forall q \in \{1, \dots N-1\}$$

The total utility π_i of u_i after all her relevant request fulfilled is the summation of all the benefits gained, subtracted by the total time cost incurred from traversing the path. Since the user may have already waited for the view requests before the current round of path planning, we add the existing values from the waiting time vector $\mathrm{w_i} = (w_{i1}, w_{i2}, \dots, w_{iN})$, in which w_{ij} denotes the time user i has already waited for the view request v_j. T_{ij} denotes the waiting time for view request v_j if the drone would travel the shortest path that visit the view requests of u_i. $\mathrm{w_i}$ and T_{ij} is updated after every round of path planning.

$$\pi_i = \sum_{s=1}^{N} g p_{is} - \sum_{j=1}^{N} w_{ij} \\ - \sum_{n=1}^{N} \sum_{m=1}^{N} \sum_{q=1}^{N-1} \sum_{k=1}^{N} p_{ik} d_{mn} (1 - \sum_{j=1}^{q} x_{kj}) x_{mq} x_{n,q+1} \tag{2}$$

The path planning problem given all view requests V can be considered as finding a feasible path $\mathrm{x} \in \mathbf{X}$ that maximizes some objective function $f(\pi_1, \pi_2, \dots, \pi_N)$ of all user utility values.

4.1.3 Optimization Objectives. Previous empirical studies [15, 16] examined people's desired resource allocation principles in situations similar to our cases, where the participant is part of a group to receive some resources but does know the exact amount that she will receive. Their results suggest that individuals prefer the principle that maximizes the average utility. Based on this finding, we choose to find the path $\mathrm{x} \in \mathbf{X}$ that maximizes the total utility of the group for Constellation. Maximizing this this objective, $\sum_{i=1}^{N} \pi_i$, effectively minimizes the total extra waiting time.

4.1.4 Algorithm. We use a branch-and-bound algorithm to search for the path in near real-time for a small number of requests (< 15). A promising future direction to improve performance is to adapt the highly efficient heuristics found for TSP [4] to our problem.

5 USER STUDY

We compared Constellation against a straightforward turn-taking baseline in a navigation task to study whether its efficiency advantage holds in the face of usability challenges.

We recruited 16 volunteers (M_{age} = 26, SD_{age} = 3.7, 8 female) from the local community to engage in two online sessions, each with eight concurrent users. Since the study focuses on interface usability, we chose an artificial target-finding task to control task difficulty. We instructed participants to find and report the labels for the landmarks in the experiment field, based on the provided descriptions.

We conducted the study remotely using a within-subject design. The participant completed the given navigation tasks in both the Constellation condition and the turn-taking condition together with seven other participants. The order of the conditions was counterbalanced between sessions. The experiment field used was a virtual theme park of world landmarks, where the 3D models of nine landmarks were laid out in a 3 by 3 grid of 50m×50m in dimension. Each landmark had one red and one blue target label on them.

From the command logs, we calculated the time between the start of the experiment and the submission of individual answers. We also calculated the average extra waiting time based on the participants' target locations, which is different from the task completion time by a constant. The Constellation condition had an apparent shorter average task completion time and extra waiting time in both sessions (Table 1). Overall, Constellation's average completion time and average extra waiting time 28.1% and 40% less than turn-taking. Constellation's evident advantage in the results demonstrated that its interface design could effectively support users in leveraging the synchronous control sharing mechanism and complete drone tours more efficiently.

6 DISCUSSION

The user utility perspective could be integrated with real-world drone path planning approaches [18, 25, 36, 53] to achieve feasible and safe paths that also accommodate diverse user inputs. Additionally, much prior research on preferred resource allocation mechanisms has focused on assigning concrete objects, such as food [28] or drinks [16]. It would be worth studying desirable drone-sharing mechanisms in a real drone tour setting in future work. Extensions to our control sharing mechanism can benefit future multi-user interfaces in other domains than drone touring. A direct application could be coordinating patients' non-critical delivery requests to be fulfilled by assistive robots in caregiving facilities. Robots could weigh the needs and locations of various items and deliver them, minimizing the waiting times.

7 CONCLUSIONS

This paper proposed Constellation, a novel multi-user drone navigation interface to serve diverse user goals in remote drone touring. Its graphical user interface allows users to flexibly communicate their interests by creating view requests or upvoting other users' requests. The underlying mechanism synthesizes user requests to generate efficient flight plans that minimize total extra waiting time. A user study with 16 participants showed that the Constellation systems could help groups navigate their points of interest more efficiently than the turn-taking alternative. In future work, we hope to develop more adaptive user models and apply similar multi-user coordination approaches to other human-agent interaction domains.

Table 1: Average task completion (TC) time and extra waiting time for Constellation and turn-taking in session 1 (s1) and 2 (s2).

	Avg TC time	Avg extra waiting time
Constellation (s1)	463s	269s
Turn-taking (s1)	617s	423s
Constellation (s2)	478s	284s
Turn-taking (s2)	692s	498s

REFERENCES

[1] Parastoo Abtahi, Benoit Landry, Jackie (Junrui) Yang, Marco Pavone, Sean Follmer, and James A. Landay. 2019. Beyond The Force: Using Quadcopters to Appropriate Objects and the Environment for Haptics in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300589

[2] Veronica Ahumada-Newhart and Judith S Olson. 2019. Going to school on a robot: Robot and user interface design features that matter. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 4 (2019), 1–28.

[3] Muhammad Zohaib Anwar, Zeeshan Kaleem, and Abbas Jamalipour. 2019. Machine learning inspired sound-based amateur drone detection for public safety applications. *IEEE Transactions on Vehicular Technology* 68, 3 (2019), 2526–2534.

[4] David L Applegate, Robert E Bixby, Vasek Chvatal, and William J Cook. 2006. *The traveling salesman problem: a computational study.* Princeton university press.

[5] Jeremy P Birnholtz, Tovi Grossman, Clarissa Mak, and Ravin Balakrishnan. 2007. An exploratory study of input configuration and group process in a negotiation task using a large display. In *Proceedings of the SIGCHI conference on Human factors in computing systems.* 91–100.

[6] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice.* Cambridge University Press.

[7] Robert S Bristow. 1999. Commentary: Virtual tourism—the ultimate ecotourism? *Tourism Geographies* 1, 2 (1999), 219–225.

[8] Anke M. Brock, Julia Chatain, Michelle Park, Tommy Fang, Martin Hachet, James A. Landay, and Jessica R. Cauchard. 2018. FlyMap: Interacting with Maps Projected from a Drone. In *Proceedings of the 7th ACM International Symposium on Pervasive Displays* (Munich, Germany) *(PerDis '18)*. Association for Computing Machinery, New York, NY, USA, Article 13, 9 pages. https://doi.org/10.1145/3205873.3205877

[9] Linfeng Chen, Kazuki Takashima, Kazuyuki Fujita, and Yoshifumi Kitamura. 2021. PinpointFly: An Egocentric Position-control Drone Interface using Mobile AR. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* 1–13.

[10] Rohan Choudhury, Gokul Swamy, Dylan Hadfield-Menell, and Anca D Dragan. 2019. On the utility of model learning in HRI. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 317–325.

[11] Erik Cohen. 1972. Toward a sociology of international tourism. *Social research* (1972), 164–182.

[12] Mark M Davis and Thomas E Vollmann. 1990. A framework for relating waiting time and customer satisfaction in a service operation. *Journal of Services Marketing* (1990).

[13] DJI. 2017. *Film Like a Pro: DJI Drone "ActiveTrack" – With Video Tutorials.* Retrieved July 13, 2020 from https://store.dji.com/guides/film-like-a-pro-with-activetrack/

[14] Eitan Frachtenberg. 2019. Practical drone delivery. *Computer* 52, 12 (2019), 53–57.

[15] Norman Frohlich, Joe A Oppenheimer, and Cheryl L Eavey. 1987. Laboratory results on Rawls's distributive justice. *British Journal of Political Science* (1987), 1–21.

[16] Vael Gates, Thomas L Griffiths, and Anca D Dragan. 2020. How to Be Helpful to Multiple People at Once. *Cognitive science* 44, 6 (2020), e12841.

[17] Christoph Gebhardt, Benjamin Hepp, Tobias Nägeli, Stefan Stevšić, and Otmar Hilliges. 2016. Airways: Optimization-based planning of quadrotor trajectories according to high-level user goals. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* 2508–2519.

[18] Christoph Gebhardt, Stefan Stevšić, and Otmar Hilliges. 2018. Optimizing for aesthetically pleasing quadrotor camera motion. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–11.

[19] Ken Goldberg, Dezhen Song, and Anthony Levandowski. 2003. Collaborative teleoperation using networked spatial dynamic voting. *Proc. IEEE* 91, 3 (2003), 430–439.

[20] Antonio Gomes, Calvin Rubens, Sean Braley, and Roel Vertegaal. 2016. *BitDrones: Towards Using 3D Nanocopter Displays as Interactive Self-Levitating Programmable Matter.* Association for Computing Machinery, New York, NY, USA, 770–780. https://doi.org/10.1145/2858036.2858519

[21] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. *Advances in neural information processing systems* 29 (2016), 3909–3917.

[22] J. C. G. Higuera, A. Xu, F. Shkurti, and G. Dudek. 2012. Socially-Driven Collective Path Planning for Robot Missions. In *2012 Ninth Conference on Computer and Robot Vision.* 417–424.

[23] Kori Inkpen, Wai-ling Ho-Ching, Oliver Kuederle, Stacey D Scott, and Garth BD Shoemaker. 1999. This is fun! We're all best friends and we're all playing!: Supporting Children's Synchronous Collaboration. (1999).

[24] Stanislav Ivanov, Ulrike Gretzel, Katerina Berezina, Marianna Sigala, and Craig Webster. 2019. Progress on robotics in hospitality and tourism: a review of the literature. *Journal of Hospitality and Tourism Technology* (2019).

[25] Niels Joubert, Mike Roberts, Anh Truong, Floraine Berthouzoz, and Pat Hanrahan. 2015. An interactive tool for designing quadrotor camera shots. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–11.

[26] Ziquan Lan, Mohit Shridhar, David Hsu, and Shengdong Zhao. 2017. XPose: Reinventing User Interaction with Flying Cameras.. In *Robotics: Science and Systems.*

[27] Walter S Lasecki, Kyle I Murray, Samuel White, Robert C Miller, and Jeffrey P Bigham. 2011. Real-time crowd control of existing interfaces. In *Proceedings of the 24th annual ACM symposium on User interface software and technology.* 23–32.

[28] Min Kyung Lee, Anuraag Jain, Hea Jin Cha, Shashank Ojha, and Daniel Kusbit. 2019. Procedural Justice in Algorithmic Fairness: Leveraging Transparency and Outcome Control for Fair Algorithmic Mediation. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 182 (Nov. 2019), 26 pages. https://doi.org/10.1145/3359284

[29] Pascal Lessel, Michael Mauderer, Christian Wolff, and Antonio Krüger. 2017. Let's Play My Way: Investigating Audience Influence in User-Generated Gaming Live-Streams. In *Proceedings of the 2017 ACM International Conference on Interactive Experiences for TV and Online Video* (Hilversum, The Netherlands) *(TVX '17)*. Association for Computing Machinery, New York, NY, USA, 51–63. https://doi.org/10.1145/3077548.3077556

[30] Jiannan Li, Ravin Balakrishnan, and Tovi Grossman. 2020. StarHopper: A Touch Interface for Remote Object-Centric Drone Navigation. In *Proceedings of Graphics Interface 2020* (University of Toronto) *(GI 2020)*. Canadian Human-Computer Communications Society / Socieétеé canadienne du dialogue humain-machine, 317 – 326. https://doi.org/10.20380/GI2020.32

[31] Andrew J. Lohn. 2017. *What's the Buzz? The City-Scale Impacts of Drone Delivery.* RAND Corporation, Santa Monica, CA. https://doi.org/10.7249/RR1718

[32] Zhicong Lu, Michelle Annett, and Daniel Wigdor. 2019. Vicariously experiencing it all without going outside: A study of outdoor livestreaming in China. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–28.

[33] David Mirk and Helmut Hlavacs. 2014. Using drones for virtual tourism. In *International Conference on Intelligent Technologies for Interactive Entertainment.* Springer, 144–147.

[34] David Mirk and Helmut Hlavacs. 2015. Virtual tourism with drones: experiments and lag compensation. In *Proceedings of the First Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use.* 45–50.

[35] Paolo Mura, Rokhshad Tavakoli, and Saeed Pahlevan Sharif. 2017. 'Authentic but not too much': exploring perceptions of authenticity of virtual tourism. *Information Technology & Tourism* 17, 2 (2017), 145–159.

[36] Tobias Nägeli, Lukas Meier, Alexander Domahidi, Javier Alonso-Mora, and Otmar Hilliges. 2017. Real-Time Planning for Automated Multi-View Drone Cinematography. *ACM Trans. Graph.* 36, 4, Article 132 (July 2017), 10 pages. https://doi.org/10.1145/3072959.3073712

[37] Yeonju Oh, Ramviyas Parasuraman, Tim McGraw, and Byung-Cheol Min. 2018. 360 vr based robot teleoperation interface for virtual tour. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI).*

[38] George Palaigeorgiou, George Malandrakis, and Christine Tsolopani. 2017. Learning with Drones: flying windows for classroom virtual field trips. In *2017 IEEE 17th International Conference on Advanced Learning Technologies (ICALT)*. IEEE, 338–342.

[39] Anastasios Plioutsias, Nektarios Karanikas, and Maria Mikela Chatzimihailidou. 2018. Hazard analysis and safety requirements for small drone operations: to what extent do popular drones embed safety? *Risk Analysis* 38, 3 (2018), 562–584.

[40] Ad Pruyn and Ale Smidts. 1998. Effects of waiting on the satisfaction with the service: Beyond objective time measures. *International journal of research in marketing* 15, 4 (1998), 321–334.

[41] Miro Puhek, Matej Perše, and Andrej Šorgo. 2012. Comparison between a real field trip and a virtual field trip in a nature preserve: knowledge gained in biology and ecology. *Journal of Baltic science education* 11, 2 (2012), 164.

[42] Dennis Ramirez, Jenny Saucerman, and Jeremy Dietmeier. 2014. Twitch plays pokemon: a case study in big g games. In *Proceedings of DiGRA.* 3–6.

[43] Mike Roberts and Pat Hanrahan. 2016. Generating Dynamically Feasible Trajectories for Quadrotor Cameras. *ACM Trans. Graph.* 35, 4, Article 61 (July 2016), 11 pages. https://doi.org/10.1145/2897824.2925980

[44] Elliot Salisbury, Sebastian Stein, and Sarvapali Ramchurn. 2015. Real-time opinion aggregation methods for crowd robotics. (2015).

[45] Dezhen Song and Ken Goldberg. 2003. Sharecam part 1: interface, system architecture, and implementation of a collaboratively controlled robotic webcam. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, Vol. 2. IEEE, 1080–1086.

[46] John I Spicer and J Stratford. 2001. Student perceptions of a virtual field trip to replace a real field trip. *Journal of Computer Assisted Learning* 17, 4 (2001), 345–354.

[47] John Stainfield, Peter Fisher, Bob Ford, and Michael Solem. 2000. International virtual field trips: a new direction? *Journal of Geography in Higher Education* 24, 2 (2000), 255–262.

[48] Daniel Szafir, Bilge Mutlu, and Terrence Fong. 2017. Designing planning and control interfaces to support user collaboration with flying robots. *The International Journal of Robotics Research* 36, 5-7 (2017), 514–542.
[49] Ryotaro Temma, Kazuki Takashima, Kazuyuki Fujita, Koh Sueda, and Yoshifumi Kitamura. 2019. Third-Person Piloting: Increasing Situational Awareness Using a Spatially Coupled Second Drone. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) *(UIST '19)*. Association for Computing Machinery, New York, NY, USA, 507–519. https://doi.org/10.1145/3332165.3347953
[50] Sébastien Thon, Dominique Serena-Allier, Céline Salvetat, and Françoise Lacotte. 2013. Flying a drone in a museum: An augmented-reality cultural serious game in Provence. In *2013 Digital Heritage International Congress (DigitalHeritage)*, Vol. 2. IEEE, 669–676.
[51] Gail Tom and Scott Lucey. 1997. A field study investigating the effect of waiting time on customer satisfaction. *The Journal of psychology* 131, 6 (1997), 655–660.
[52] Panos Trahanias, Wolfram Burgard, Antonis Argyros, Dirk Hahnel, Haris Baltzakis, Patrick Pfaff, and Cyrill Stachniss. 2005. TOURBOT and WebFAIR: Web-operated mobile robots for tele-presence in populated exhibitions. *IEEE Robotics & Automation Magazine* 12, 2 (2005), 77–89.
[53] Ke Xie, Hao Yang, Shengqiu Huang, Dani Lischinski, Marc Christie, Kai Xu, Minglun Gong, Daniel Cohen-Or, and Hui Huang. 2018. Creating and chaining camera moves for qadrotor videography. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.
[54] Yaxing Yao, Huichuan Xia, Yun Huang, and Yang Wang. 2017. Privacy mechanisms for drones: Perceptions of drone controllers and bystanders. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 6777–6788.