

Creating the Virtual Reality of the Future with Artificial Intelligence and Computational Design

Lap-Fai (Craig) Yu

Associate Professor
Computer Science Department
George Mason University



Design Computing and eXtended Reality

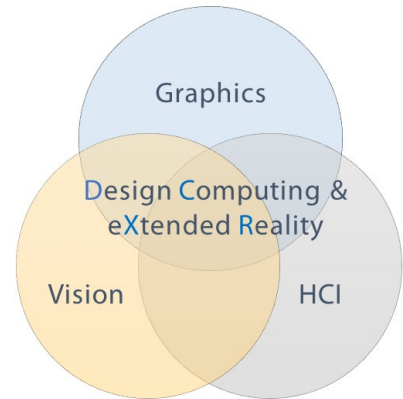
What will the future of work, entertainment, and everyday life be 5-10 years from now?



Interests: Computer Graphics, Computer Vision, HCI


Techniques: Optimization, Artificial Intelligence, Machine Learning, Simulation

Applications: VR/AR, Computational Design, User Interfaces



AI for Design (e.g., interior design, architectural design, product design)


[SIGGRAPH 2015]



[SIGGRAPH Asia 2016]




[SIGGRAPH 2016]




[SIGGRAPH Asia 2020]



Initial Layout



Target Layout




[SIGGRAPH 2011]




Scene Understanding and Reconstruction (e.g., stereo, affordance analysis)


[ICCV 2017]



[ICCV 2015]




[CVPR 2013]



[3DV 2016]



[VR 2019]



[CHI 2021]



Virtual Reality Adaptive Training (e.g., driving, disaster response, exergaming, education)

[VR 2018]



[VR 2018]




[VR 2017]



[SIGGRAPH 2020]



[VR 2020]




Computational Interaction (e.g., novel wayfinding tool, virtual experiences, synthesized sound/speech)

[CHI 2019]



[CHI 2019]




[CHI 2021]



A woman is seen speaking to the camera and leads into her playing a routine.

The crowd cheers for the people.

[CHI 2019]



[SIGGRAPH Asia 2019]





This is the United States Capitol, the home of the United States Congress.



Computational Design

Computational Design Framework



Design Solution Space



Computational Design Framework

Design Goals

- Functionality (e.g., ergonomics, physical properties)
- Aesthetics prior (e.g., color, style)



Computational Design Framework

Design Constraints



- Space
- Budget (e.g., cost, material, time, labor)



Computational Design Framework

Design Variables



- Positions of objects



Computational Design Framework

Design Variables


- Orientations of objects




Computational Design Framework

Search for Solutions

- Optimization
- Statistical Inference (e.g., MCMC)




Initialization



good solution 1



good solution 2



Low Cost

High Cost

(Good Design)

(Bad Design)

Virtual Environment Synthesis

Virtual Environment Synthesis


“Make it Home: Automatic Optimization of Furniture Arrangement”, SIGGRAPH 2011

Lap-Fai Yu, Sai-Kit Yeung, Chi-Keung Tang, Demetri Terzopoulos, Tony F. Chan, Stanley J. Osher




Design Acceleration

Traditional (CAD):




Our Method (Computational Design):



Our Approach

Learning Step


Positive Examples




Relationship
Extraction:
1) Spatial
2) Hierarchical
3) Pairwise




Initialization



Optimization Step



Synthesis Finished




Optimization: Total Cost Function

$$C(\phi) = w_a C_a(\phi) + w_v C_v(\phi) + w_{\text{path}} C_{\text{path}}(\phi) \quad \text{Ergonomics}$$
$$+ w_{\text{pr}}^d C_{\text{pr}}^d(\phi) + w_{\text{pr}}^\theta C_{\text{pr}}^\theta(\phi) \quad \left. \vphantom{C(\phi)} \right\} \text{Prior}$$
$$+ w_{\text{pair}}^d C_{\text{pair}}^d(\phi) + w_{\text{pair}}^\theta C_{\text{pair}}^\theta(\phi)$$

$$\phi = \{(p_i, \theta_i) \mid i = 1 \dots n\}$$

Position, Orientation



Optimization

Simulated Annealing

- Computational imitation of physical annealing process

Cooling schedule:


- At the beginning, high temperature:
 - “heat up” furniture objects, allow flexible rearrangement
- Over time, temperature lowers gradually:
 - rearrangement is less aggressive
- At the end, temperature drops to zero:
 - refine final arrangement



Optimization

- At each iteration, a “move” is proposed,
 - Moves: translation, rotation, swapping objects, moving pathway controls

- Transition:



- Metropolis criterion determines **transition probability**:

$$\alpha(\phi' | \phi) = \min \left[\frac{f(\phi')}{f(\phi)}, 1 \right]$$

$$= \min \left[\exp(\beta(C(\phi) - C(\phi'))), 1 \right]$$

,where $f(\phi) = e^{-\beta C(\phi)}$

C decreases, $\alpha = 1$

C increases, $0 < \alpha < 1$

$$\beta \propto \frac{1}{\text{Temperature}}$$

Design Exploration




3D Walkthrough



Architectural Layout Synthesis

“**Crowd-driven Mid-scale Layout Design**”, SIGGRAPH 2016


Tian Feng, Lap-Fai Yu, Sai-Kit Yeung, KangKang Yin, Kun Zhou



Synthesis



Architectural Layout Design




Input

- Layout domain
- Possible sites
(e.g., boutiques, restaurants, restrooms)

Output

- Floor plan

Architectural Layout Design



Input

- Layout domain
- Possible sites
(e.g., boutiques, restaurants, restrooms)

Design Goals?



[Incorporating Human Factors](#)

Output

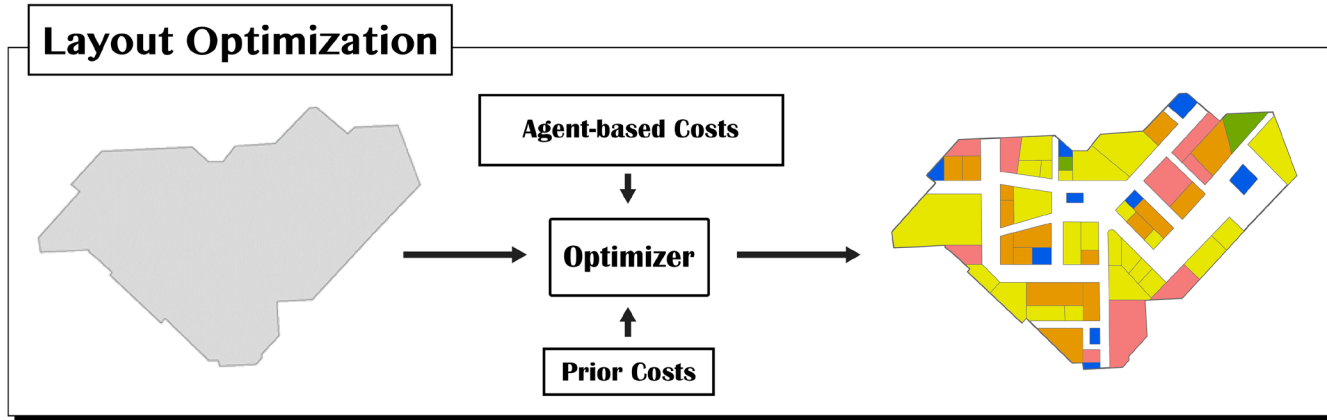
- Floor plan

Crowd Simulation

- Typical Metric: Crowd Density
- Traditional: Evaluation
- Ours: Evaluation + Optimization



Layout Optimization



$$C(\phi) = \mathbf{C}_A \mathbf{w}_A^T + \mathbf{C}_P \mathbf{w}_P^T$$


\mathbf{C}_A : Agent-based Cost
(Mobility, Accessibility, Coziness)

Computed by agent-based simulation

\mathbf{C}_P : Prior Cost
(Floor Area Ratio, Total # of sites, # of each type of sites)

Computed by real world layouts' statistics

Mobility: how smooth is the agents' walking experience



Low Cost




High Cost

actual speed


$$C_m(\phi) = 1 - \frac{1}{N} \sum_i \frac{\bar{v}_i}{d_i}$$

baseline speed

Accessibility: how reasonably sites are distributed



Low Cost



High Cost

walking distance between two sites

$$C_a(\phi) = \frac{1}{NL} \sum_i \frac{1}{k_i} \sum_j l_{i,j}$$

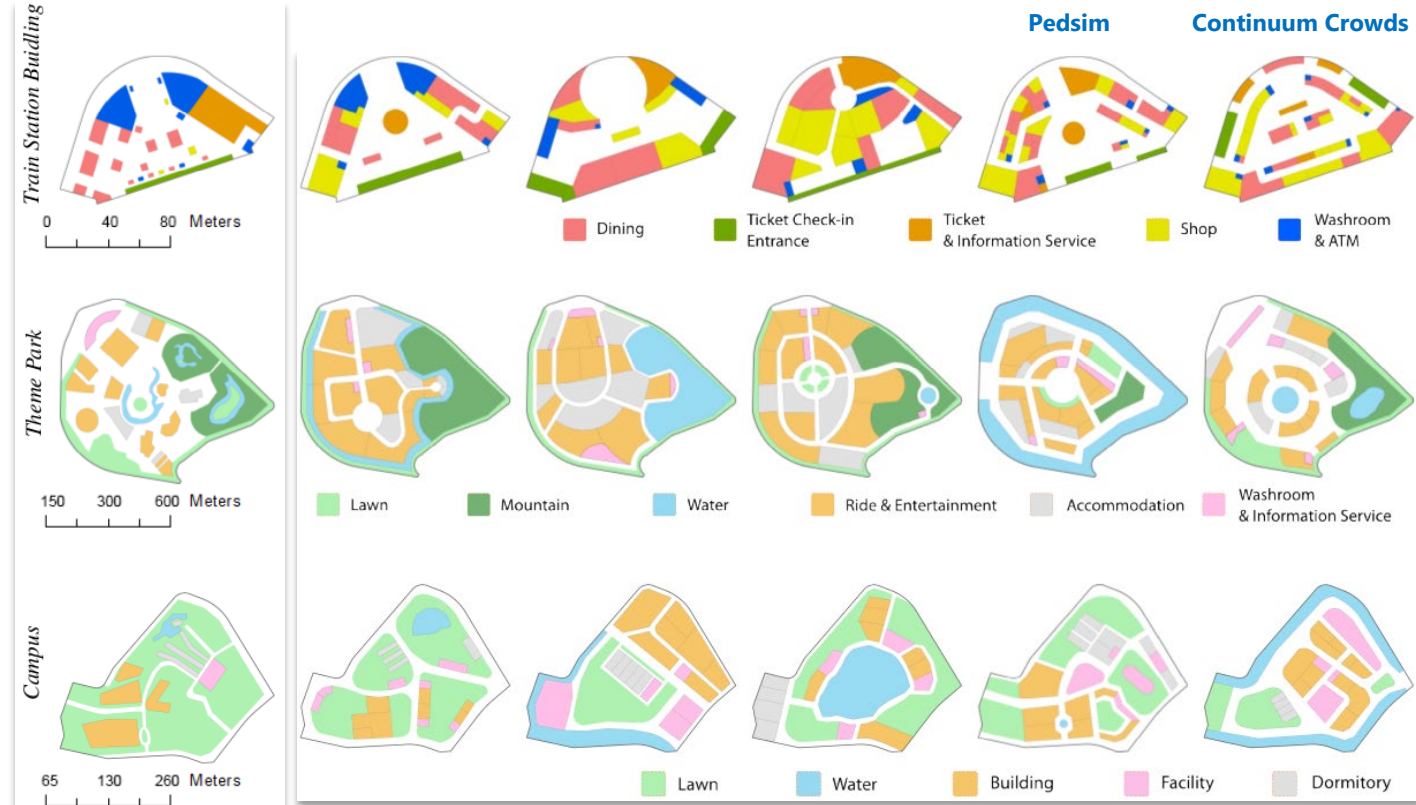
Coziness: how good agents feel in sites



crowd density in the site being visited

$$C_c(\phi) = \frac{1}{N} \sum_i \frac{1}{q_i} \sum_j \left[1 - \exp\left(-\frac{(\rho_{i,j}^{\text{visit}} - \mu)^2}{2\sigma^2}\right) \right]$$

Results





Real World Layout

Syntheses

Mall 2



Real World Layout

Synthesis







Real World Layout

Synthesis




Train Station Building



Perceptual Data-Guided Computational Design

- Human user data tracked in **Virtual Reality**
- Head, hand, body, eye gaze, EEG, etc.



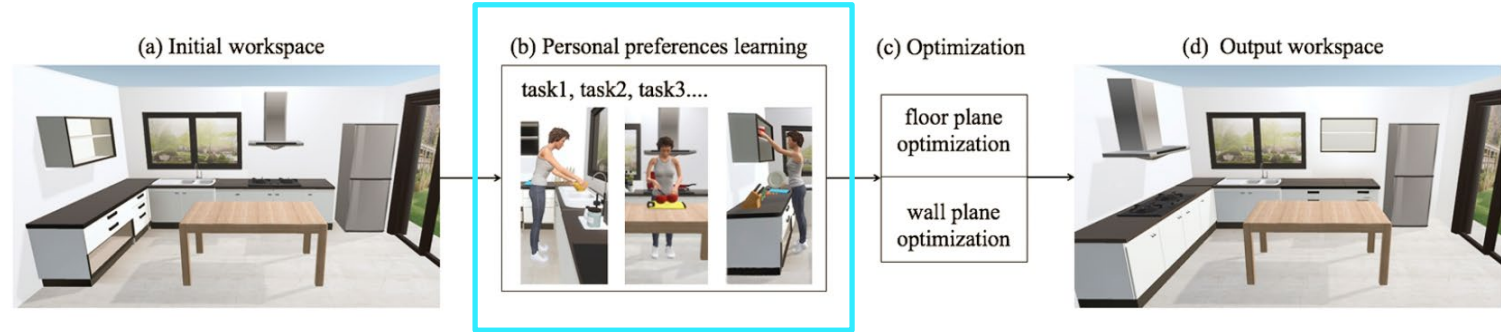
[Courtesy of Stephan Lu]

Learning Workspace Preferences from VR

“Functional Workspace Optimization via Learning Personal Preferences from Virtual Experiences”, *IEEE VR, 2019*

Wei Liang, Jingjing Liu, Yining Lang, Bing Ning, Lap-Fai Yu





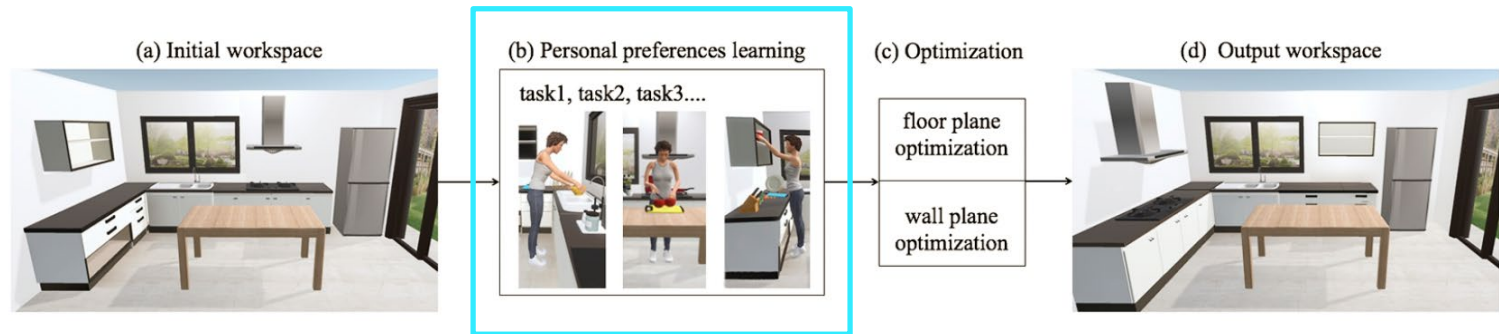


Fig. 3: An example of personal preferences. Each colored bar depicts the user interacting with the component with the same color in the scene. The ordered sequence of the components represents the personal preferences for the workflow.



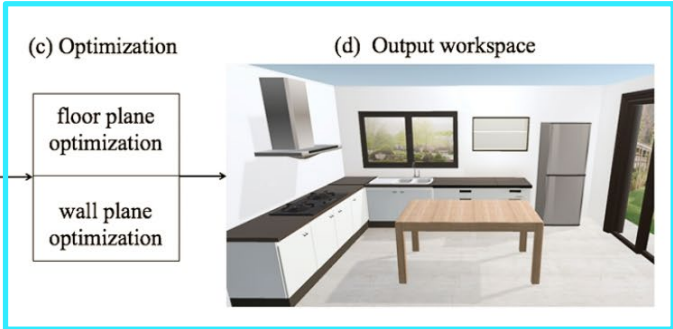



Fig. 4: The position distribution when the user interacts with the corresponding component. We use a heatmap to visualize the visited times at each point when the user interacts with the corresponding component. The redder the color is, the more times the user visited that point.



Personalized VR Driving Training

“Synthesizing Personalized Training Programs for Improving Driving Habits via Virtual Reality”, *IEEE VR, 2018*

Yining Lang, Wei Liang, Fang Xu, Yibiao Zhao, Lap-Fai Yu

Eye-tracking VR headset



Pedestrians:




Turning:




Changing Lane:



Front Car:





Driving Habits:

Signal before a Turn




Signal before Changing Lane




Stop for Pedestrians




Optimized Training Routes:



(a) More Turns



(b) Fewer Turns



(c) Pass Specified Positions

VR Training:

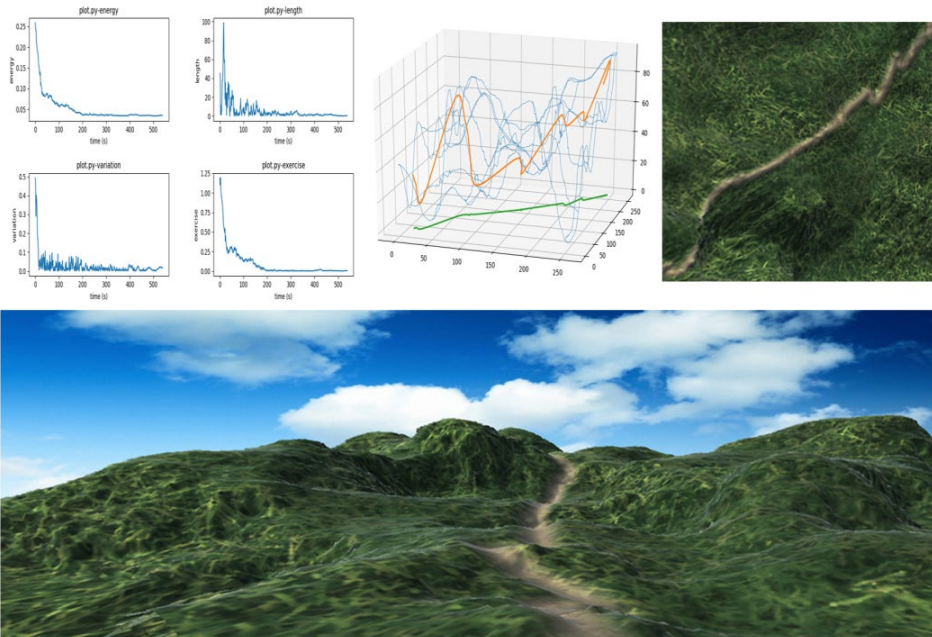


Gaze tracked in VR



Terrain Generation for VR Biking

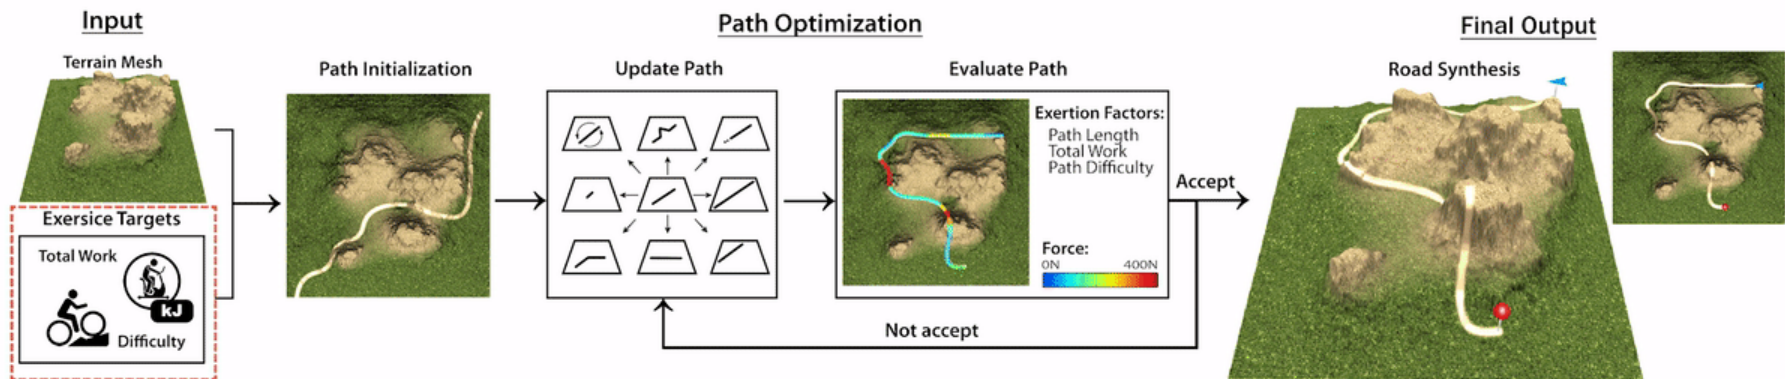
Fitbit VR bike:




“Exertion-Aware Path Generation”, SIGGRAPH 2020

Wanwan Li, Biao Xie, Yongqi Zhang, Walter Meiss, Haikun Huang, Lap-Fai Yu

Overview





VR Exercise Biking





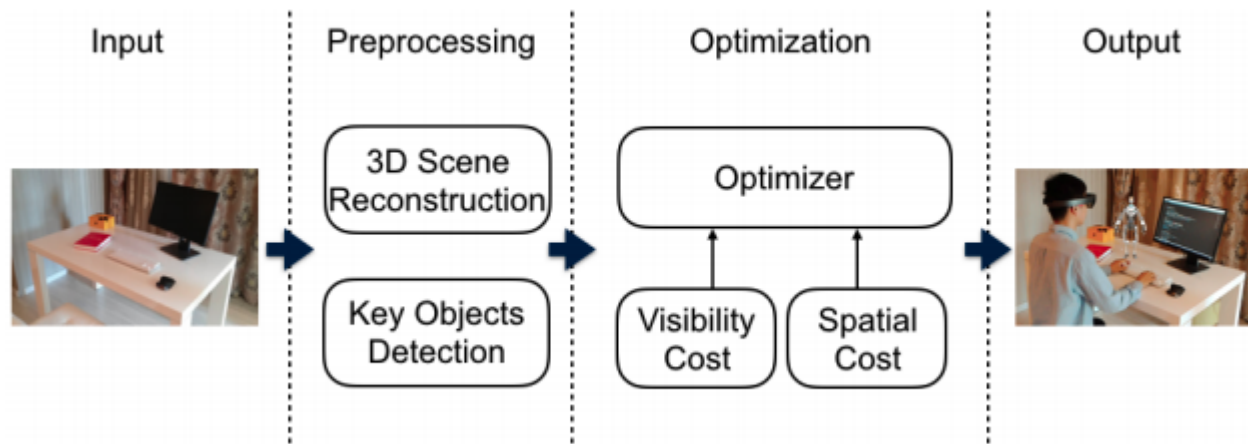
Scene-Aware Virtual Agents

Virtual Agent Positioning

“Virtual Agent Positioning Driven by Scene Semantics in Mixed Reality”, *IEEE VR*, 2019

Yining Lang, Wei Liang, Lap-Fai Yu






Virtual Pet Synthesis

“Scene-Aware Behavior Synthesis for Virtual Pets in Mixed Reality”, ACM CHI 2021

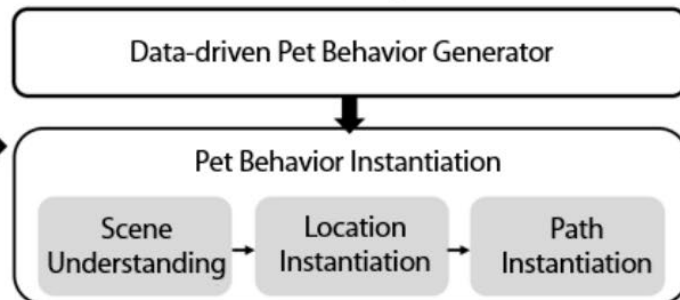
Wei Liang, Xinzhe Yu, Rawan Alghofaili, Yining Lang, Lap-Fai Yu




(a) Input



(b) Pet Behavior Synthesis



(c) Output



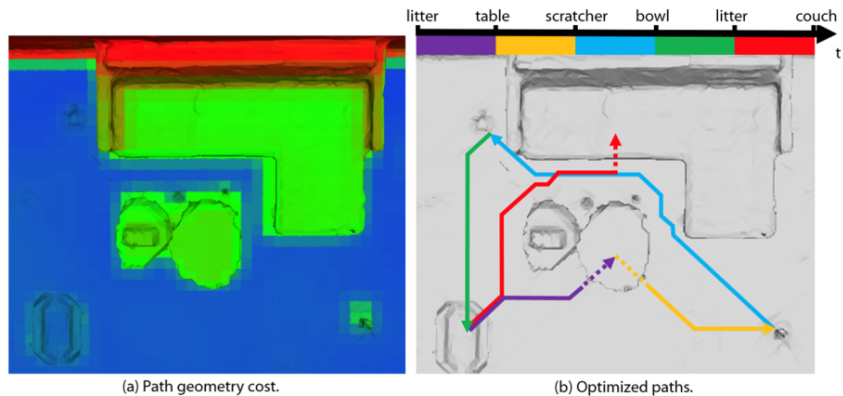


Fig. 6. (a) A visualization of the path geometry cost for each cell. The redder a cell is, the higher its cost value is. (b) An illustration of the optimized paths. The state bar at the top shows the sequence of objects that the pet travels to. Each color refers to a path going from one object's location to another object's location (e.g., purple refers to going from the litter to the table).

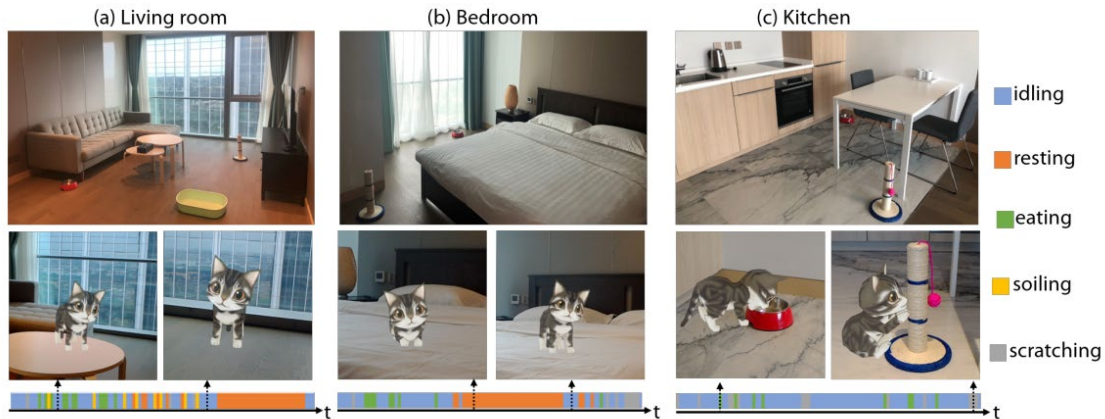


Fig. 7. The living room, bedroom, and kitchen scenes used in our experiments. The top row shows the input scene. The middle row demonstrates some generated behaviors. The bottom row shows the generated behavior sequence visualized by a state bar. Each sequence consists of 100 behaviors and each behavior is shown by one color in the bar.


Navigation Aid

Adaptive VR Navigation Aid



“Lost in Style: Gaze-Driven Adaptive Aid for VR Navigation”, *CHI 2019*


Rawan Alghofaili, Yasuhito Sawahata, Haikun Huang, Hsueh-Cheng Wang, Takaaki Shiratori, Lap-Fai Yu




Navigation Aids in Virtual Worlds





Arrow



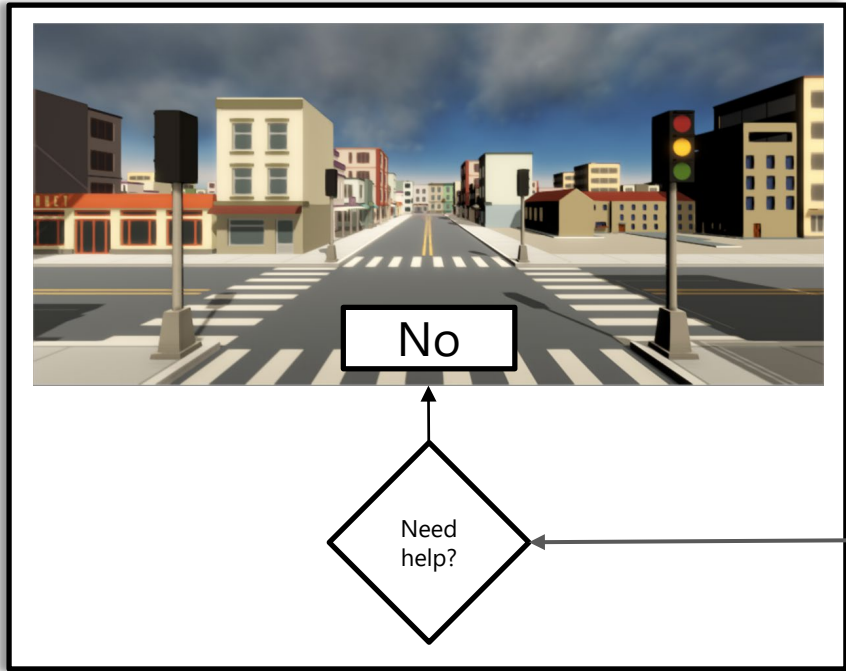
Mini Map




FOVE VR eye-tracking headset

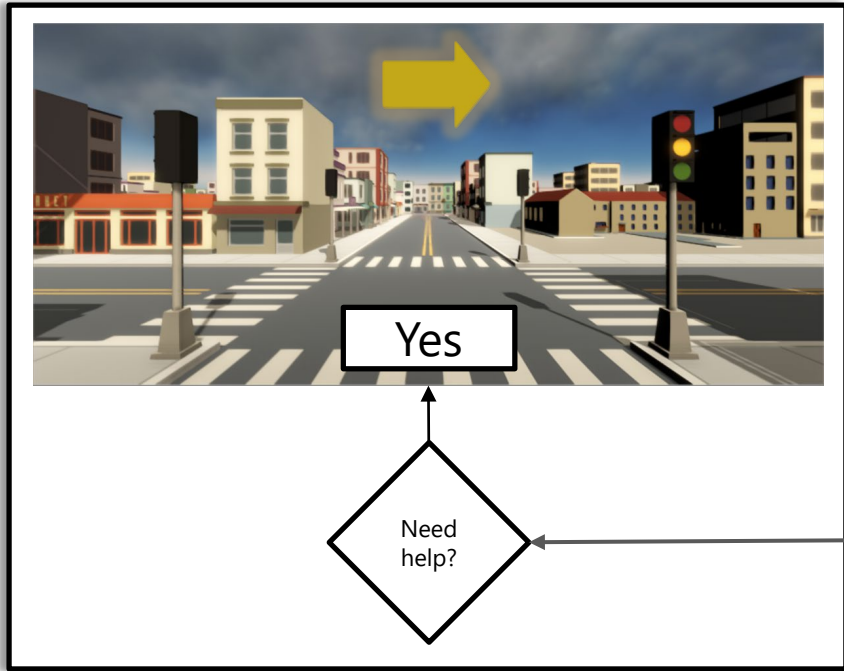


Eye-tracking gaze sequences




Navigation Need Classifier (LSTM)






Navigation Need Classifier (LSTM)

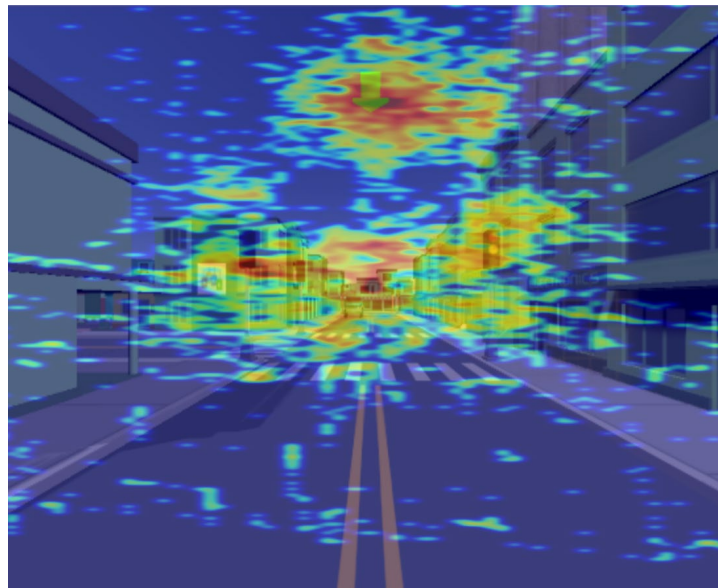




Adaptive navigation arrow



Permanent navigation arrow



Dawn of VR/AR gaming

- What will be the de facto tools / techniques?


Game designer's end

content creation / authoring
level design
animation
gameplay / genres


Player's end

interaction
user interface
collaboration
communication





4 Reals of VR



The screen is a window through which one sees a virtual world. The challenge is to make that world look real, act real, sound real, feel real.


— *Ivan Sutherland* —

AZ QUOTES


Real yet?

Act real?

AR Agent Positioning [VR 19]




Affordance Reasoning [ICCV 17]




Sound real?

Audible Panorama [CHI 19]



Feel real?

Tactile inference? Haptics-aware 3D synthesis?




Shared realism


Multi-user VR Experiences



PhD students:



Haikun Huang



Noah Duncan



Tian Feng




Mahdi Jampour



Yujia Wang



Rawan Ghofaili



Changyang Li



Wanwan Li

MS students:



Michael Solah



Yining Lang



Monica Lin

Acknowledgements: Students & Mentees

Undergraduate students:



Carla Aravena



Darian Springer



Biao Xie



Hanqing Wang



Lorenzo Barrett



Yongqi Zhang



Mark Vo



Javier Talavera



Amilcar Samayoa



Siem Sium

Acknowledgements

- Industry Collaborators:
 - Adobe Research, Microsoft Research, Facebook Oculus, Wayfair
- Funding:
 - NSF CRII, "CHS: Perceptual Data-Guided Computational Design"
 - NSF CAREER, "Performance-Guided Synthesis of Virtual Environments for Personalized Training"
 - NSF Graduate Research Fellowship
 - Adobe Research
- McNair Program
- Louis Stokes Alliances for Minority Participation (LSAMP)
- GMU OSCAR Program



Please check out my research website
for many projects and demos!

<https://cs.gmu.edu/~craigyu/>