Delicate Dance:

Preferences in Interactive Meshing

Changkun Ou
changkun.de/s/polyred5star

@changkun

IDC 2021 Spring
Virtual Event
Munich
Mar 15, 2021

Factual Error

Bad English

Lack of Explanation

Code Syntax Error

Tons of Typos

Incomprehensible

Wrong Information

Previously on Polygon Reduction (Polyred)...

Previously on Polygon Reduction (Polyred)...

A Glimpse into Advances of
Mesh Representation Learning

Changkun Ou
https://changkun.de/s/polyred1step

IDC 2019 Spring
Bernried, Germany
Apr 3, 2019

changkun.de/s/polyred1step (2019a)

Simplicity is Complicated:
On the Balance of Performance and Knobs

Changkun Ou
https://changkun.de/s/polyred2what

IDC 2019 Autumn
Vienna, Austria
Oct 9, 2019

changkun.de/s/polyred2what (2019b)

Polygon Reduction:
Under the Hood

Changkun Ou
https://changkun.de/s/polyred3

IDC 2020 Spring
Bernried, Germany
Mar 31, 2020

~~changkun.de/s/polyred3 (2020a)~~
404 PANDEMIC

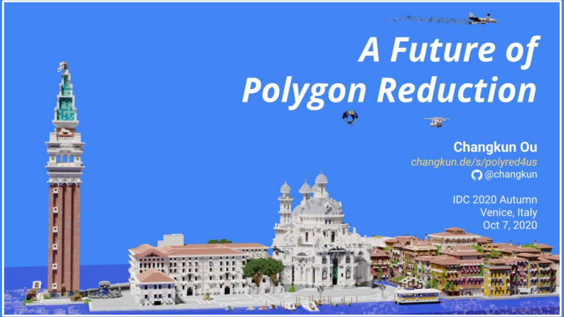

changkun.de/s/polyred4us (2020b)

Previously on Polygon Reduction (Polyred)...

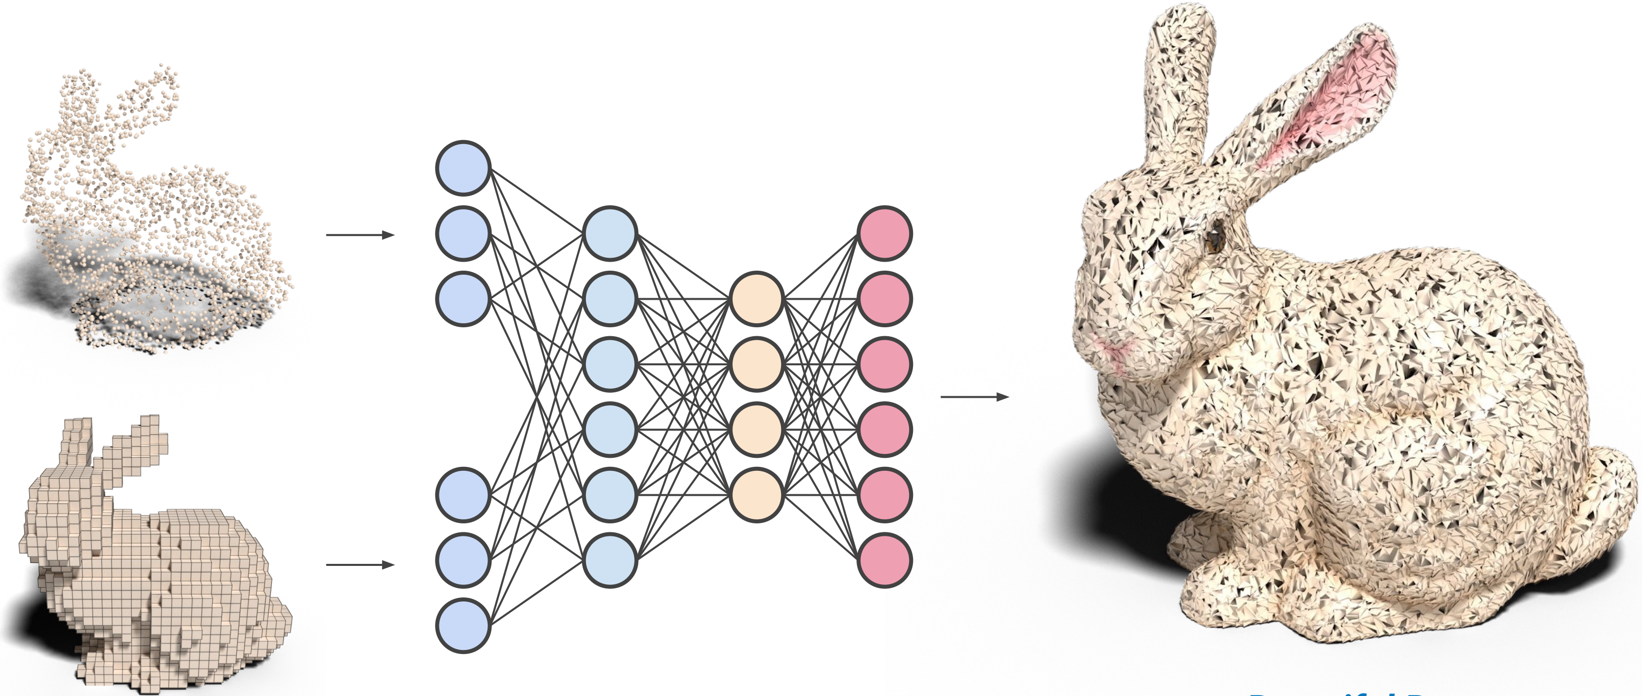

Beautiful But Terrible Meshing

Previously on Polygon Reduction (Polyred)...

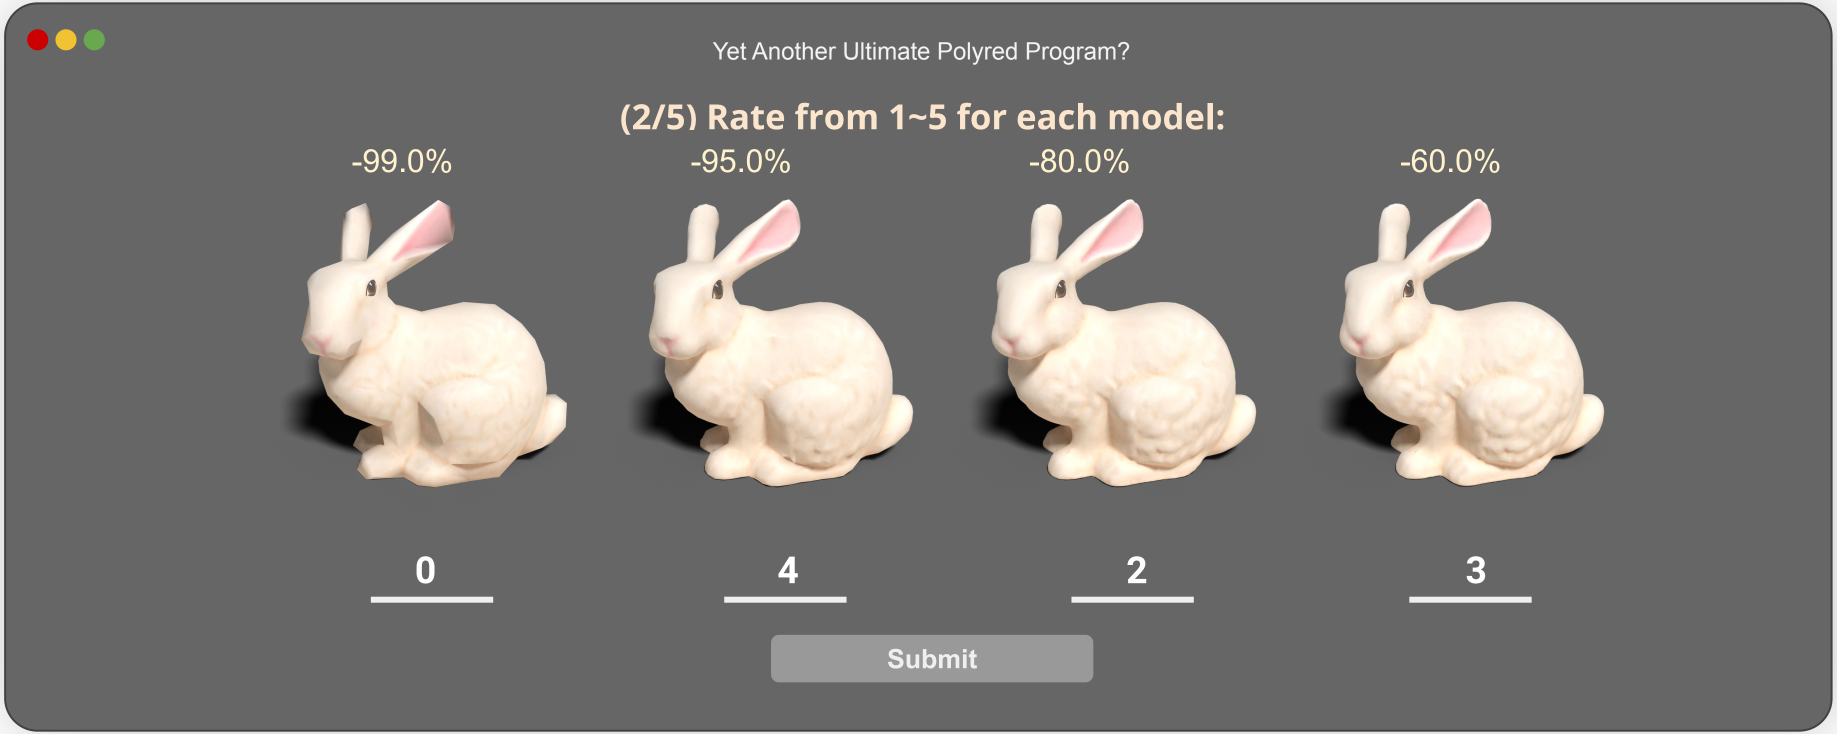

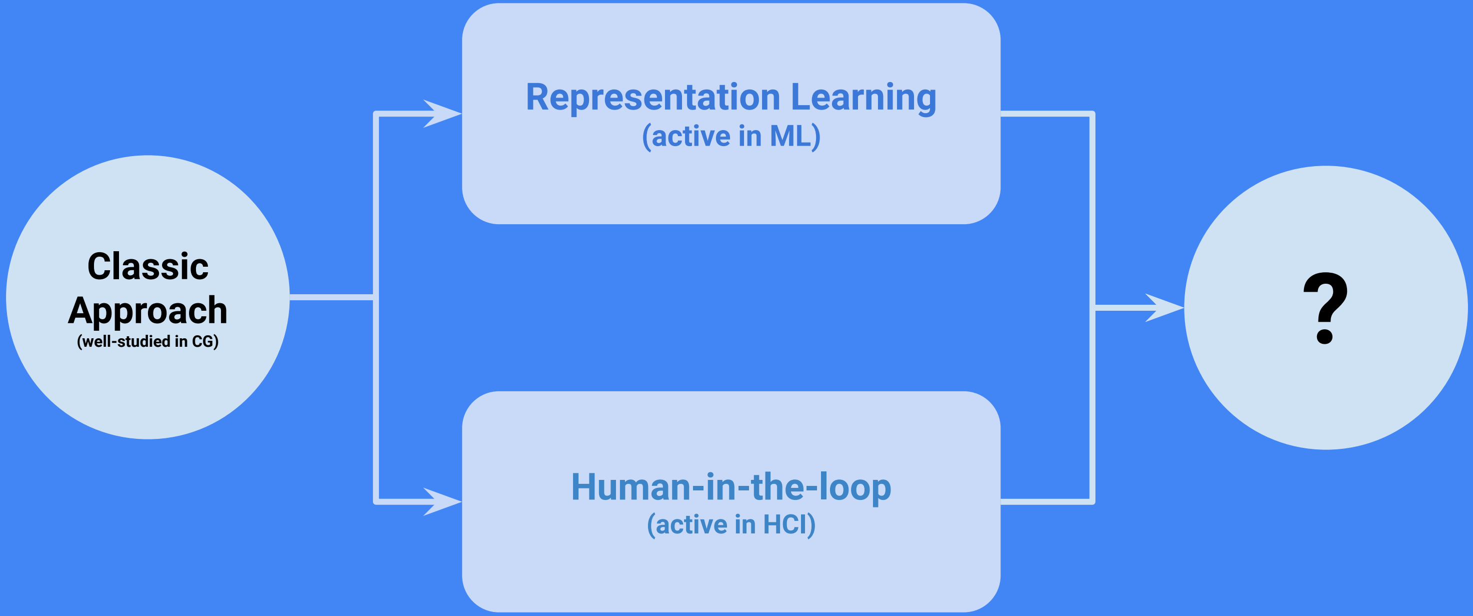
Classic Approach
(well-studied in CG)
Representation Learning
(active in ML)
Human-in-the-loop
(active in HCI)
?

(So what?) "I think theses works are between communities, which community to you want to contribute? (Pick a side)"

-- Prof. Schmidt

Essentials

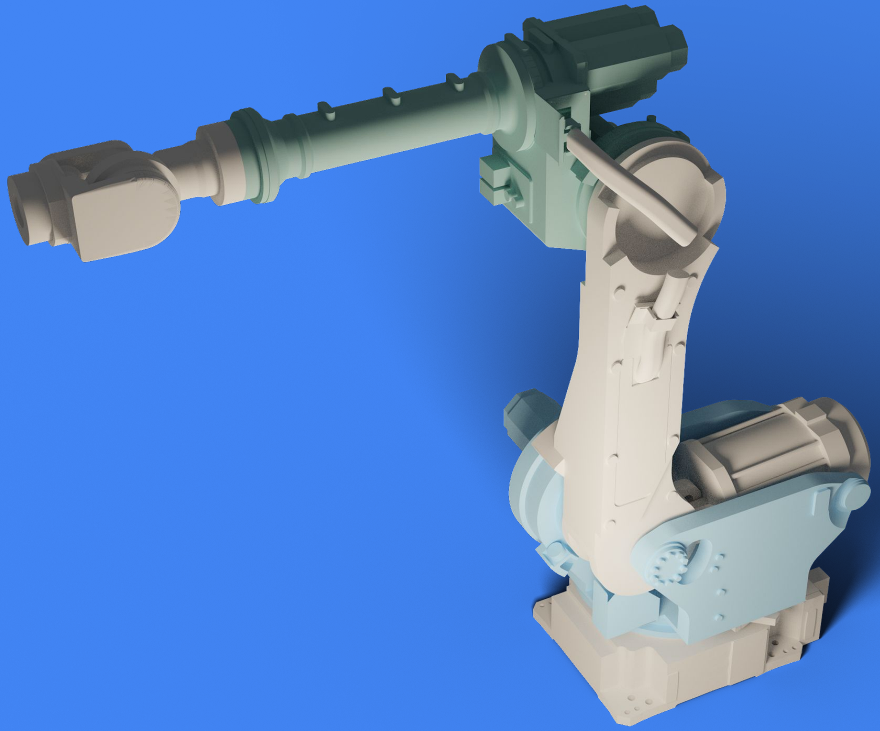

Design Task as Optimization Problem

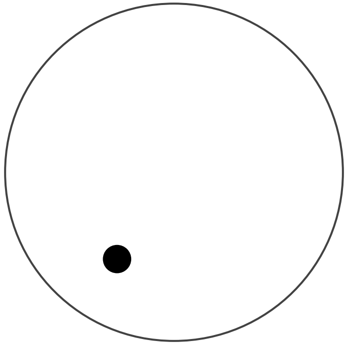

design space

Design Task as Optimization Problem

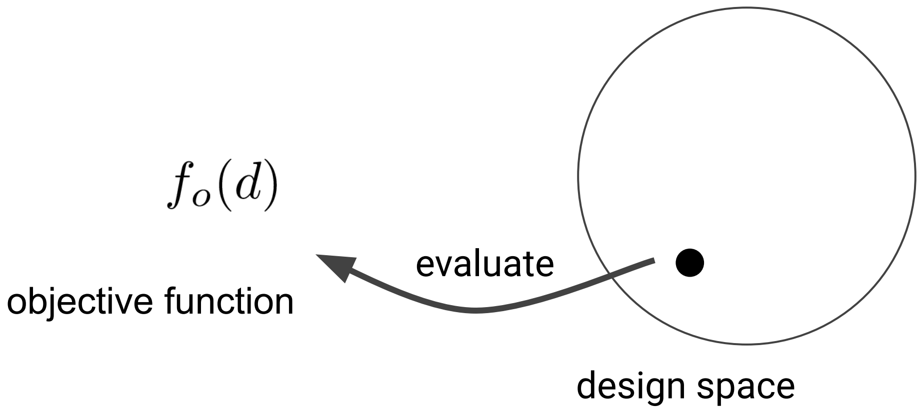

Design Task as Optimization Problem

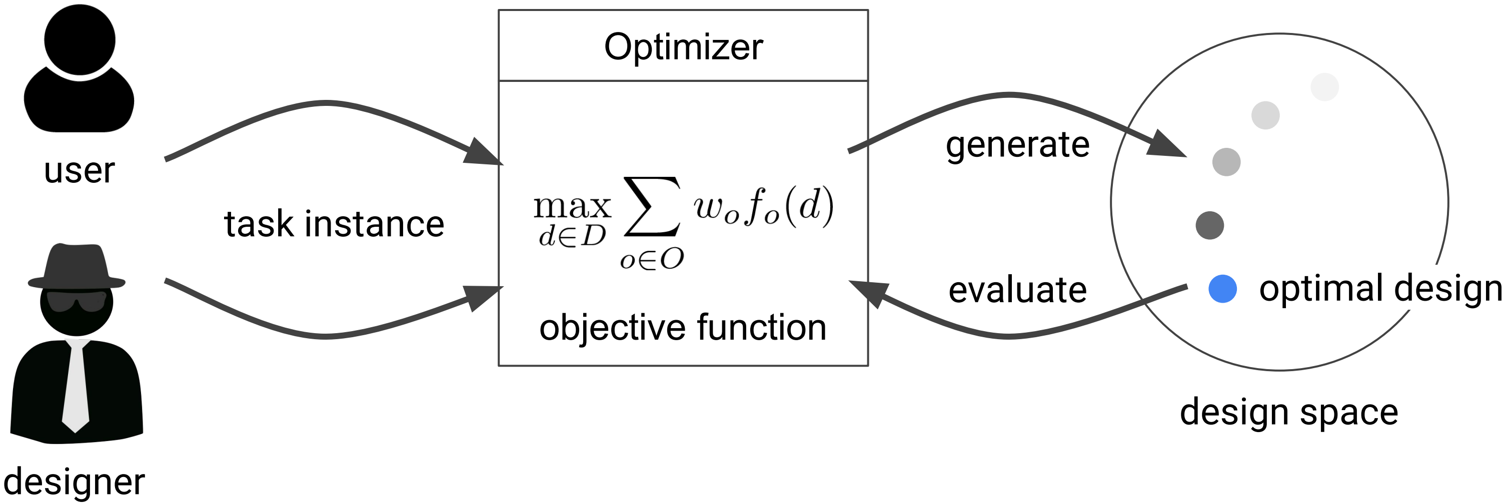

- See successful works in keyboard layout optimization by A. Oulasvirta, D. Buschek, etc...

Instantiation: 3D Modeling as Optimization Problem

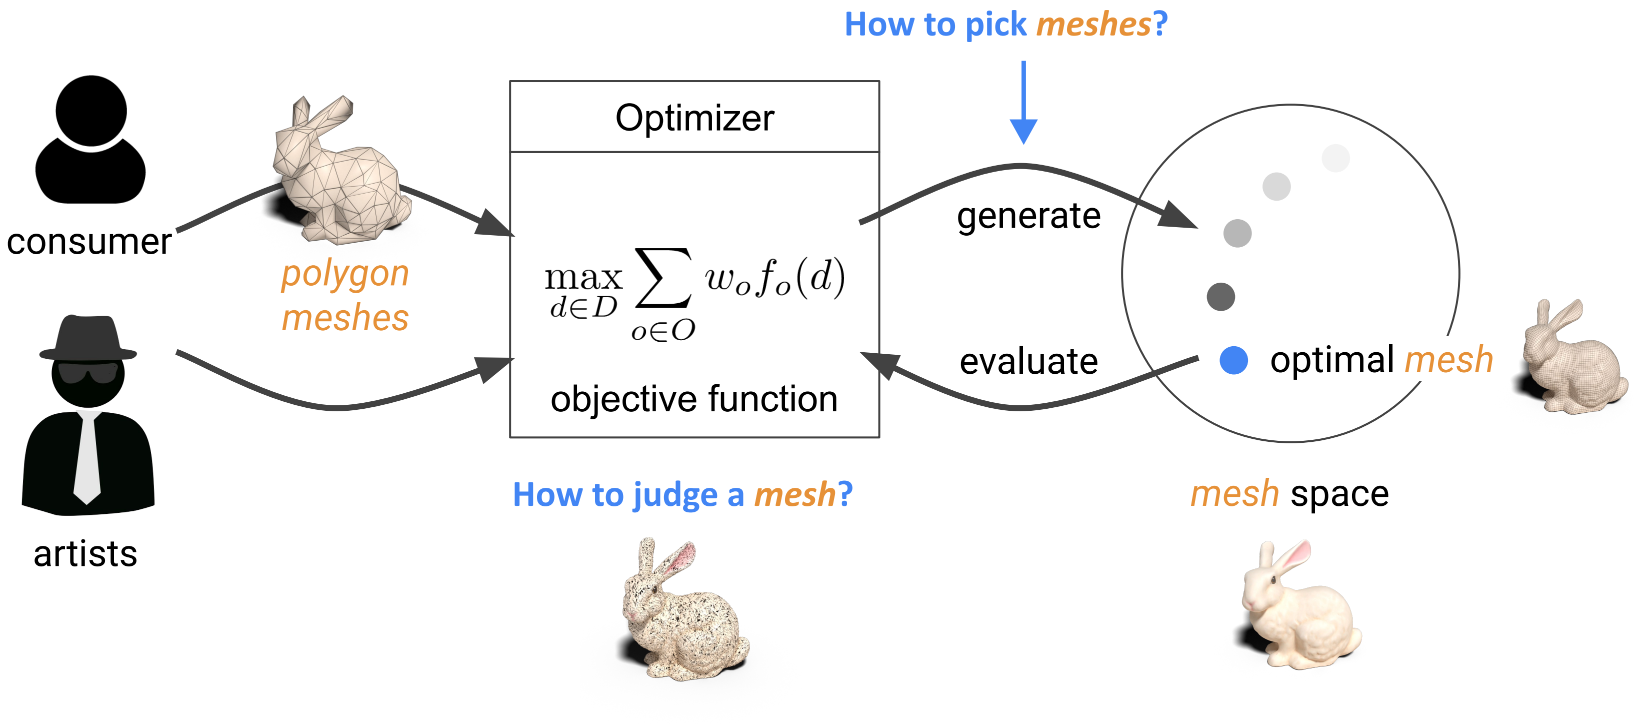

Time Complexity on Interactive 3D Modeling

Total Time Spent

Decision Time to Pick a Design

Single-step Modeling Time

$$T = \sum_{a \in A} \left(DT(\mathbf{x}) + P(\mathbf{m}) \right)$$

#Attempts

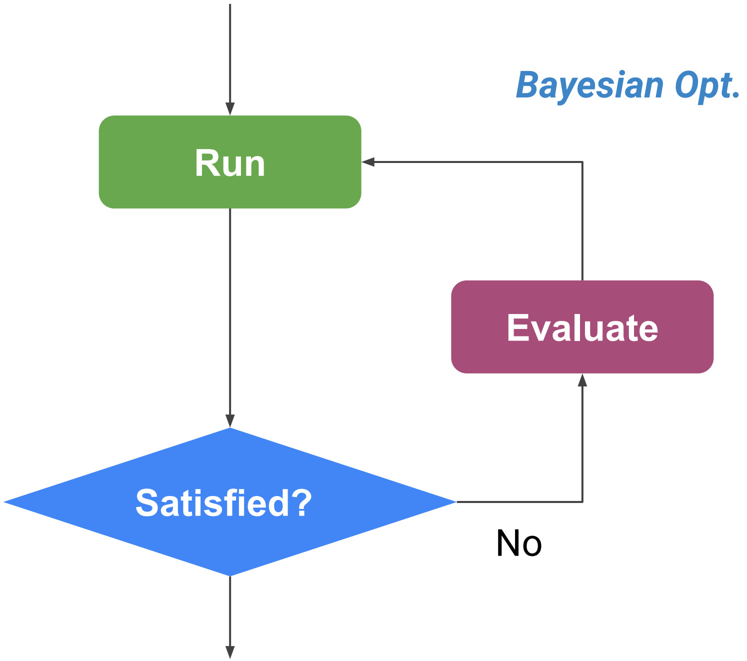

Time Complexity on Interactive 3D Modeling

Total Time Spent — Decision Time to Pick a Design — Single-step Modeling Time

$$T = \sum_{a \in A} (DT(\mathbf{x}) + P(\mathbf{m}))$$

#Attempts

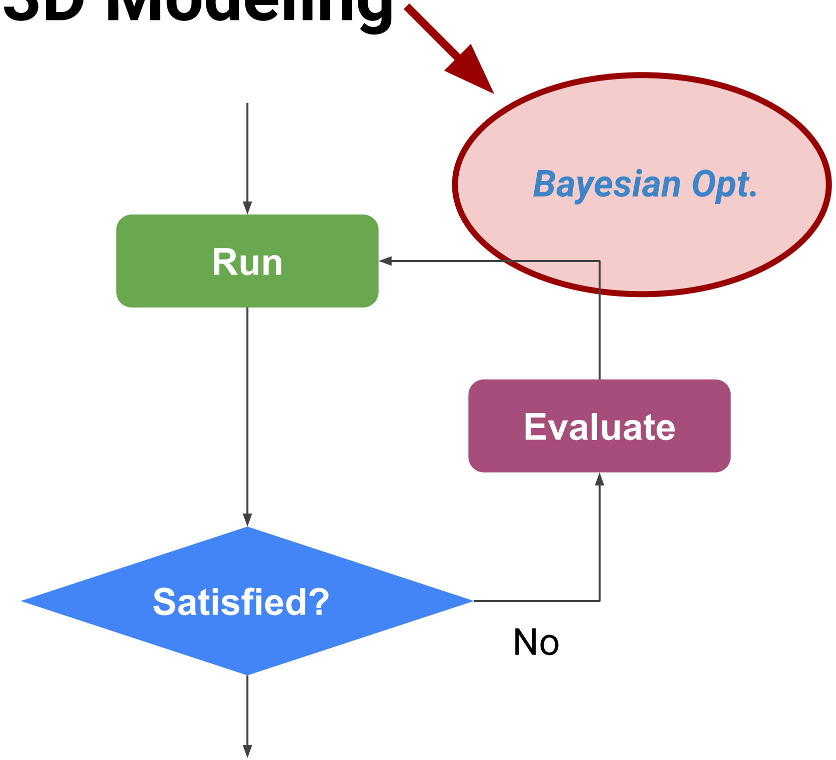

Design Optimization Problem Specification

- Design space $\mathcal{X}$ (or say search space) is parameterized by $x \in \mathcal{X}$
- The goal is to search the best parameter settings s.t.

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmax}}\, g(\mathbf{x})$$

- where g is the user's preference (and expensive to evaluate)
- Bayesian optimization is widely used for hyperparameter search with *few* queries: $\mathbf{x}^* \succ \{\mathbf{x}_{(i)}\}_{i=1}^{m}$
- Bring reinforcement learning agent strategies
  - *Exploration* (i.e. How to propose the next parameter settings?)
  - *Exploitation* (i.e. How to converge as fast as possible?)

A computer graphics artist sits down to use a simple renderer to find appropriate surfaces for a typical reflectance model. It has a series of parameters that must be set to control the simulation: “specularity”, “Fresnel reflectance coefficient”, and other, less-comprehensible ones. The parameters interact in ways difficult to discern. The artist knows in his mind’s eye what he wants, but he’s not a mathematician or a physicist — no course he took during his MFA covered Fresnel reflectance models. Even if it had, would it help? He moves the specularity slider and waits for the image to be generated. The surface is too shiny. He moves the slider back a bit and runs the simulation again. Better. The surface is now appropriately dull, but too dark. He moves a slider down. Now it’s the right colour, but the specularity doesn’t look quite right any more. He repeatedly bumps the specularity back up, rerunning the renderer at each attempt until it looks right. Good. Now, how to make it look metallic...?

Geometry Galleries

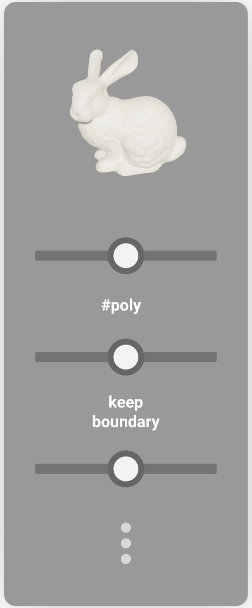

Geometry Galleries

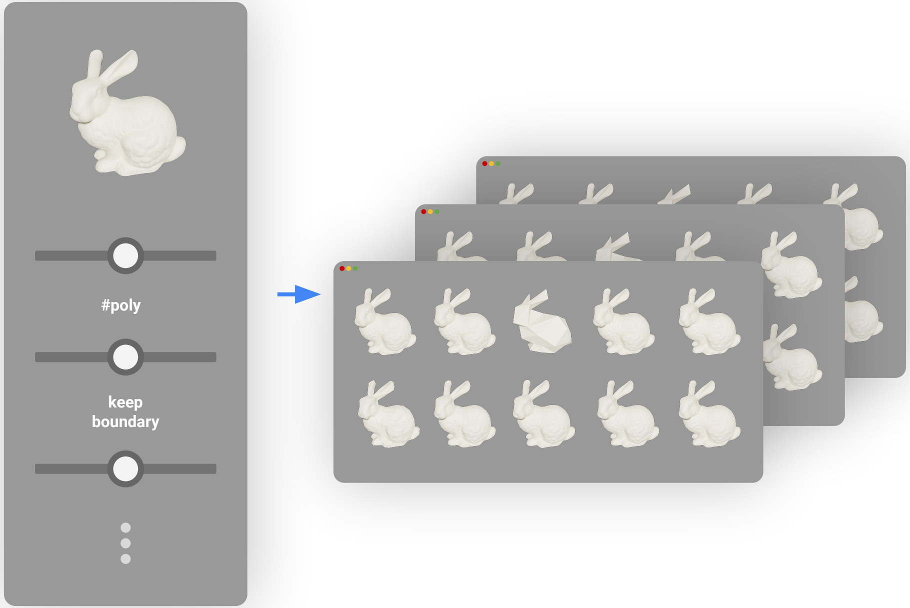

Geometry Galleries

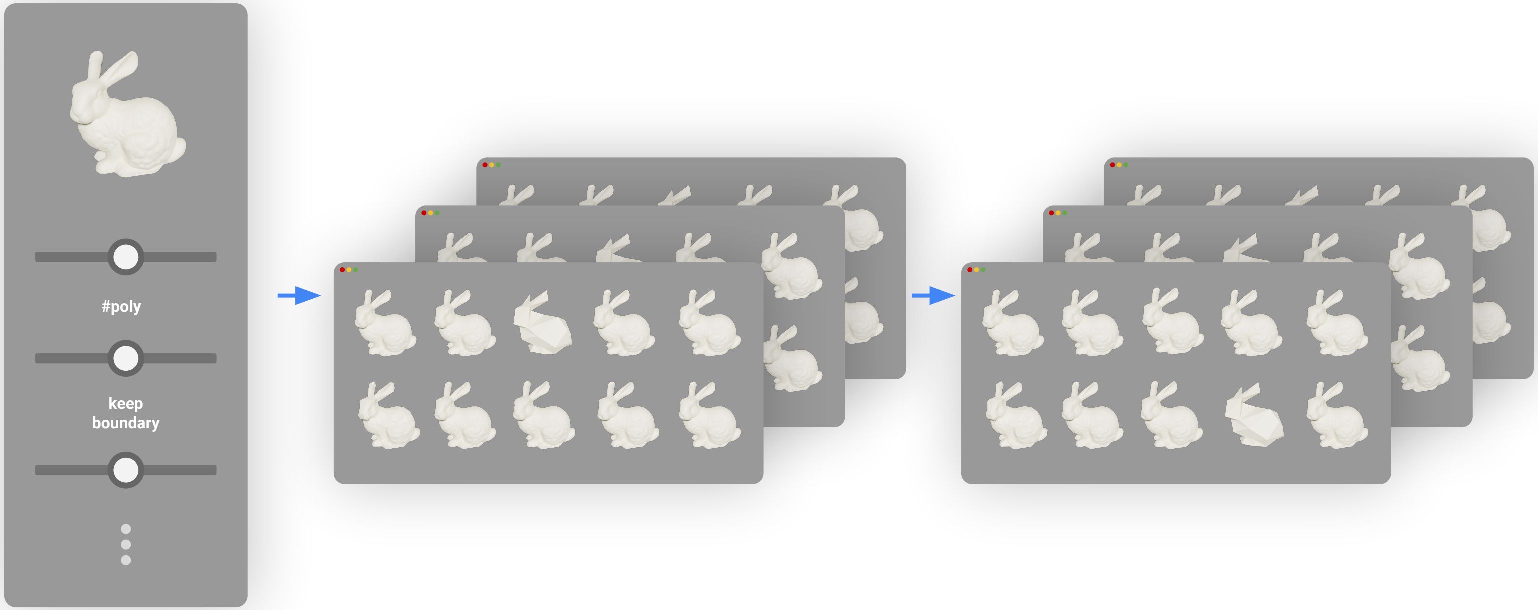

Geometry Galleries

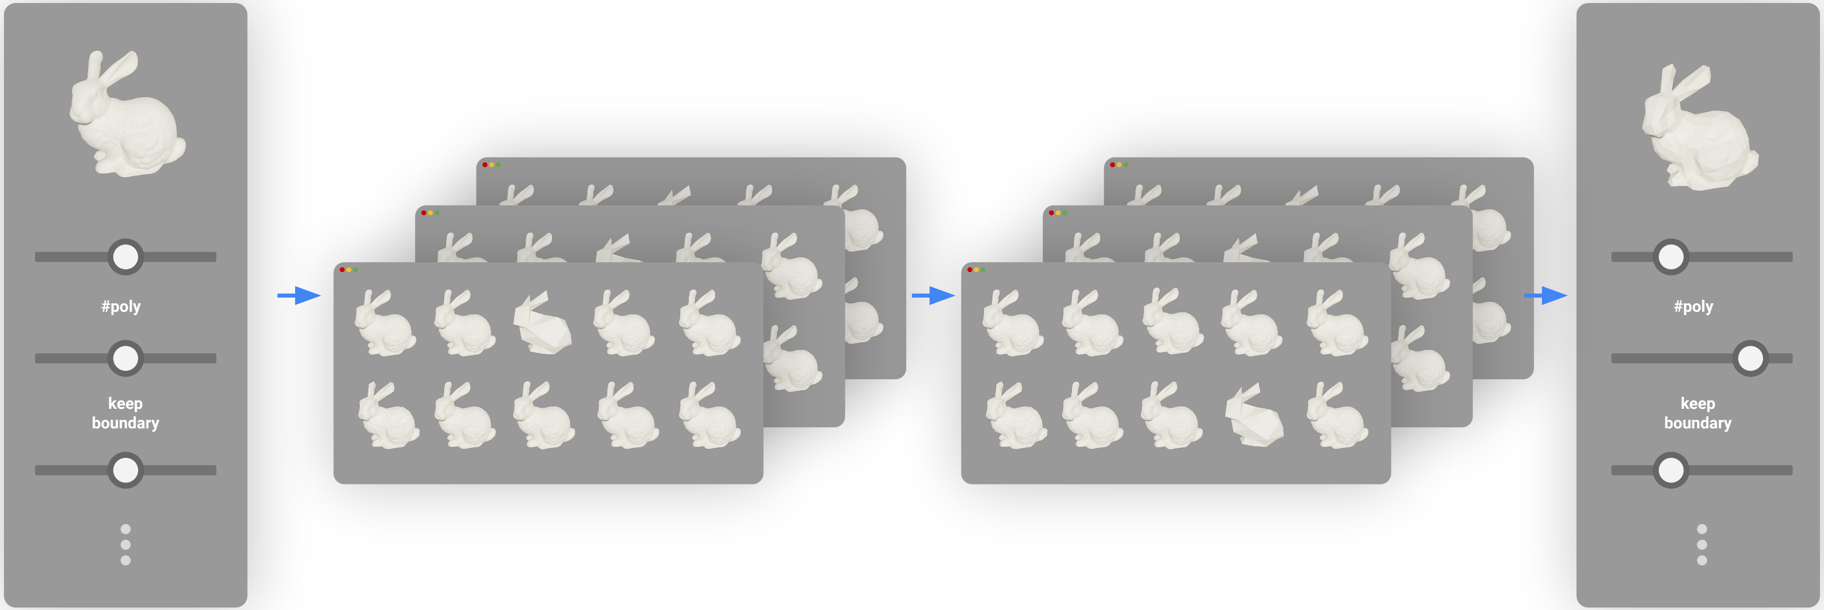

What are the actual preferences among these meshes?

- (Almost) Equal polygon counts
- Preference record from a 3D designer

4 4 4 4

What are the actual preferences among these meshes?

- Exactly equally reduced polygon counts
- Preference record from another senior 3D designer

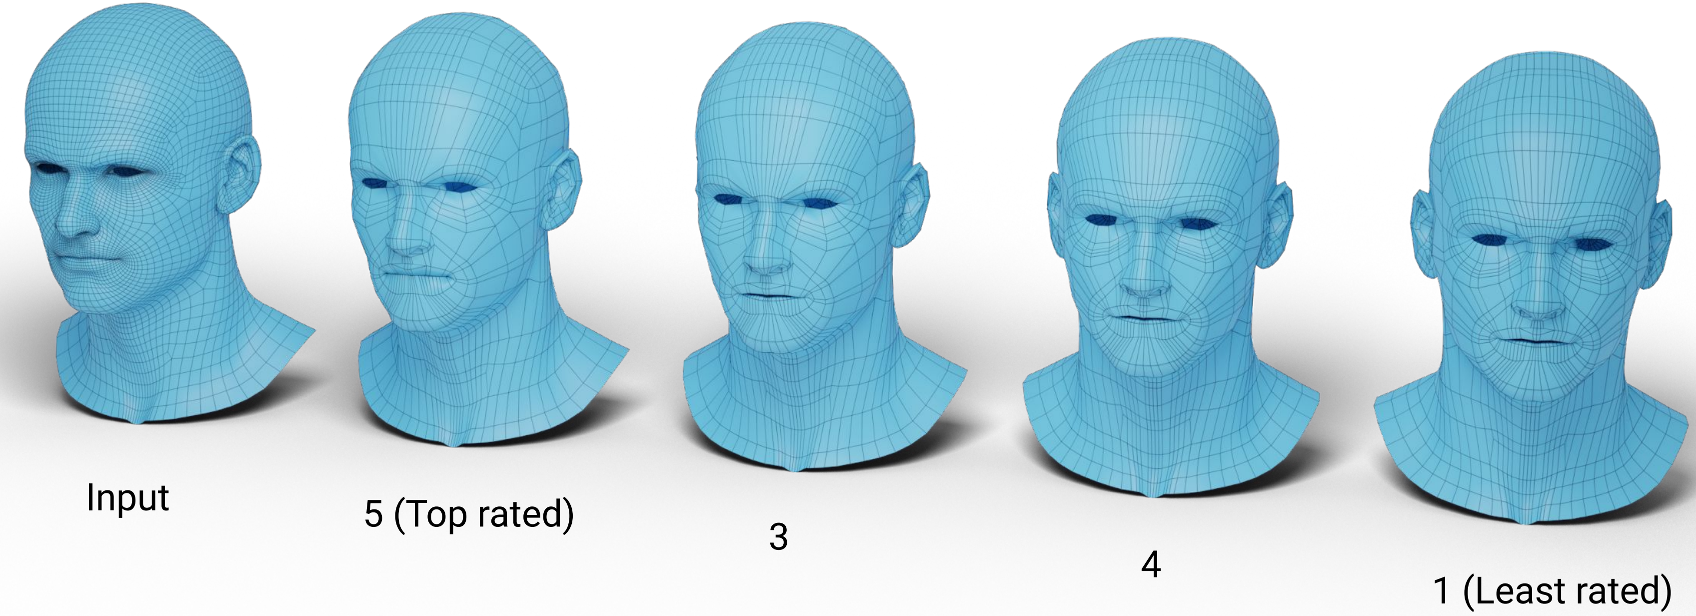

A Further Comparison

1 v.s. 4 (zero difference, just *z-fighting*)

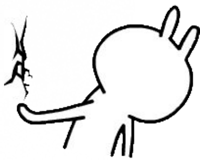

Why?

- Maybe your experiment was not designed properly…
- Maybe your algorithm's output space have no overlap to the desired design space…
- Maybe your participants is not reliable…
- Maybe …

I don't know! - These results commonly exist in real world cases

Theories

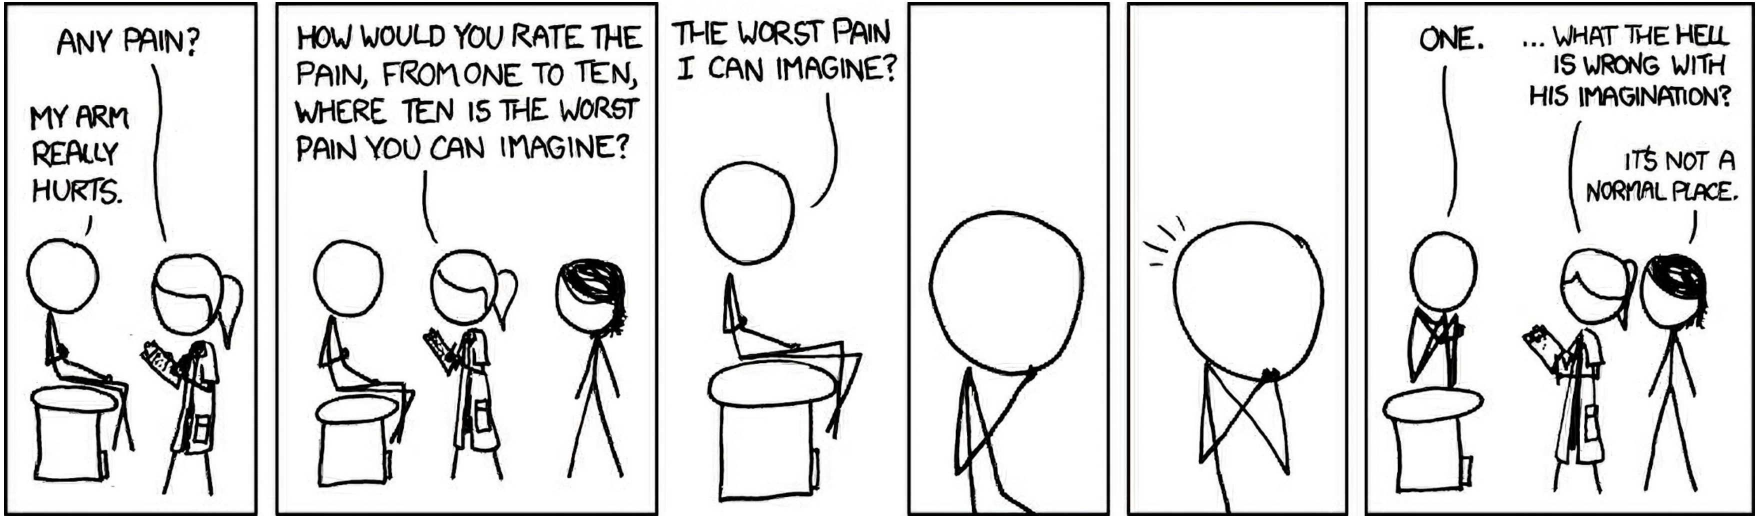

***Pain Rating** at xckd.com/883*

"My arm really hurts."

"How would you rate the pain, from 1 to 10, where 10 is the worst pain you can imagine?"

"One."

Judgement under Uncertainty **by A. Tversky and D. Kahneman**

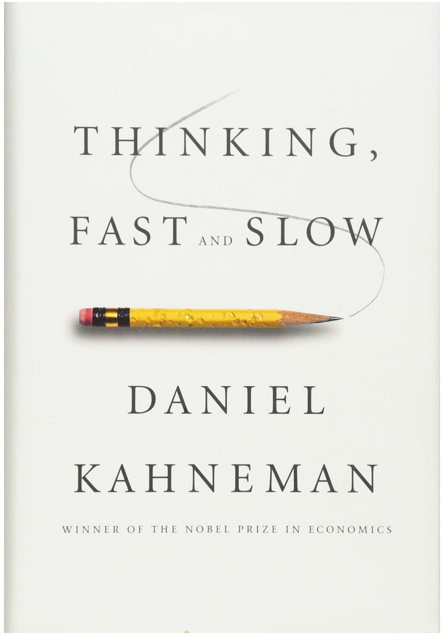

Economic Choices†

By DANIEL MCFADDEN*

This Nobel lecture discusses the microeconometric analysis of choice behavior of consumers who face discrete economic alternatives. Before the 1960's, economists used consumer theory mostly as a logical tool, to explore conceptually the properties of alternative market organizations and economic policies. When the theory was applied empirically, it was to market-level or national-accounts-level data. In these applications, the theory was usually developed in terms of a *representative agent*, with market-level behavior given by the representative agent's behavior writ large. When observations deviated from those implied by the representative agent theory, these differences were swept into an additive disturbance and attributed to data measurement errors, rather than to unobserved factors within or across individual agents. In statistical language, traditional consumer theory placed structural restrictions on mean behavior, but the distribution of responses about their mean was not tied to the theory.

In the 1960's, rapidly increasing availability of survey data on individual behavior, and the advent of digital computers that could analyze these data, focused attention on the variations in demand across individuals. It became important to explain and model these variations as part of consumer theory, rather than as ad hoc disturbances. This was particularly obvious for discrete choices, such as transportation mode or occupation. The solution to this problem has led to the tools we have today for microeconometric analysis of choice behavior. I will first give a brief history of the development of this subject, and place my own contributions in context. After that, I will discuss in some detail more recent developments in the economic theory of choice, and modifications to this theory that are being forced by experimental evidence from cognitive psychology. I will close with a survey of statistical methods that have developed as part of the research program on economic choice behavior.

Science is a cooperative enterprise, and my work on choice behavior reflects not only my own ideas, but the results of exchange and collaboration with many other scholars.[1] First, of course, is my co-laureate James Heckman, who among his many contributions pioneered the important area of dynamic discrete choice analysis. Nine other individuals who played a major role in channeling microeconometrics and choice theory toward their modern forms, and had a particularly important influence on my own work, are Zvi Griliches, L. L. Thurstone, Jacob Marschak, Duncan Luce, Amos Tversky, Danny Kahneman, Moshe Ben-Akiva, Charles Manski, and Kenneth Train. A gallery of their photographs is shown in Figure 1. I wish particularly to cite Griliches, Marschak, and Tversky, robbed by death of their own chances to win Nobel prizes.

I. A Brief History

Classical economic theory postulates that consumers seek to maximize their self-interest, and that self-interest has broadly defined consistency properties across different decisions. At one level, the theory is virtually tautological, as in this description from a principles textbook by Frank Taussig (1912):

> An object can have no value unless it has utility. No one will give anything for an article unless it yield him satisfaction. Doubtless people are sometimes foolish, and buy things, as children do, to please a moment's fancy; but at least they think at

† This article is a revised version of the lecture Daniel McFadden delivered in Stockholm, Sweden on December 8, 2000, when he received the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel. The article is copyright © The Nobel Foundation 2000 and is published here with the permission of the Nobel Foundation.

* Department of Economics, University of California, Berkeley, CA 94720. Many of the author's publications cited in this paper are posted at http://elsa.berkeley.edu/~mcfadden.

[1] Any accounting of credit for my contributions to economics has to include Leo Hurwicz, John Chipman, Marc Nerlove, and Hirofumi Uzawa, who attracted me to the field and taught me most of what I know.

351

Comparative Judgement **by L. Thurstone, R. Luce, R. Bradley**

$$P(\text{choose A over B})$$

Psychological Review
1994, Vol. 101, No. 2, 266–270

A Law of Comparative Judgment

L. L. Thurstone

The object of this paper is to describe a new psychophysical law which may be called the *law of comparative judgment* and to show some of its special applications in the measurement of psychological values. The law of comparative judgment is implied in Weber's law and in Fechner's law. The law of comparative judgment is applicable not only to the comparison of physical stimulus intensities but also to qualitative comparative judgments such as those of excellence of specimens in an educational scale and it has been applied in the measurement of such psychological values as a series of opinions on disputed public issues. The latter application of the law will be illustrated in a forthcoming study. It should be possible also to verify it on comparative judgments which involve simultaneous and successive contrast.

The law has been derived in a previous article and the present study is mainly a description of some of its applications. Since several new concepts are involved in the formulation of the law it has been necessary to invent several terms to describe them, and these will be repeated here.

Let us suppose that we are confronted with a series of stimuli or specimens such as a series of gray values, cylindrical weights, handwriting specimens, children's drawings, or any other series of stimuli that are subject to comparison. The first requirement is of course a specification as to what it is that we are to judge or compare. It may be gray values, or weights, or excellence, or any other quantitative or qualitative attribute about which we can think 'more' or 'less' for each specimen. This attribute which may be assigned, as it were, in differing amounts to each specimen defines what we shall call the *psychological continuum* for that particular project in measurement.

As we inspect two or more specimens for the task of comparison there must be some kind of process in us by which we react differently to the several specimens, by which we identify the several degrees of excellence or weight or gray value in the specimens. You may suit your own predilections in calling this process psychical, neural, chemical, or electrical but it will be called here in a non-committal way *the discriminal process* because its ultimate nature does not concern the formulation of the law of comparative judgment. If then, one handwriting specimen *seems* to be more excellent than a second specimen, then the two discriminal processes of the observer are different, at least on this occasion.

Editor's Note. This article is a reprint of an original work published in 1927 in the *Psychological Review, 34,* 273–286. Comments by R. Duncan Luce and Robyn M. Dawes follow this article.

L. L. Thurstone, University of Chicago.

This article was one of a series of articles by members of the Behavior Research Staff of the Illinois Institute for Juvenile Research, Chicago, *Herman M. Adler,* Director. Series B, No. 107.

The so-called 'just noticeable difference' is contingent on the fact that an observer is not consistent in his comparative judgments from one occasion to the next. He gives different comparative judgments on successive occasions about the same pair of stimuli. Hence we conclude that the discriminal process corresponding to a given stimulus is not fixed. It fluctuates. For any handwriting specimen, for example, there is one discriminal process that is experienced more often with that specimen than other processes which correspond to higher or lower degrees of excellence. This most common process is called here *the modal discriminal process for the given stimulus.*

The psychological continuum or scale is so constructed or defined that the frequencies of the respective discriminal processes for any given stimulus form a normal distribution on the psychological scale. This involves no assumption of a normal distribution or of anything else. The psychological scale is at best an artificial construct. If it has any physical reality we certainly have not the remotest idea what it may be like. We do not assume, therefore, that the distribution of discriminal processes is normal on the scale because that would imply that the scale is there already. We *define* the scale in terms of the frequencies of the discriminal processes for any stimulus. This artificial construct, the psychological scale, is so spaced off that the frequencies of the discriminal processes for any given stimulus form a normal distribution on the scale. The separation on the scale between the discriminal process for a given stimulus on any particular occasion and the modal discriminal process for that stimulus we shall call *the discriminal deviation* on that occasion. If on a particular occasion, the observer perceives more than the usual degree of excellence or weight in the specimen in question, the discriminal deviation is at that instant positive. In a similar manner the discriminal deviation at another moment will be negative.

The standard deviation of the distribution of discriminal processes on the scale for a particular specimen will be called its *discriminal dispersion.*

This is the central concept in the present analysis. An ambiguous stimulus which is observed at widely different degrees of excellence or weight or gray value on different occasions will have of course a large discriminal dispersion. Some other stimulus or specimen which is provocative of relatively slight fluctuations in discriminal processes will have, similarly, a small discriminal dispersion.

The scale difference between the discriminal processes of two specimens which are involved in the same judgment will be called *the discriminal difference* on that occasion. If the two stimuli be denoted A and B and if the discriminal processes corresponding to them be denoted a and b on any one occasion, then the discriminal difference will be the scale distance $(a - b)$ which varies of course on different occasions. If, in one of the comparative judgments, A seems to be better than B, then, on

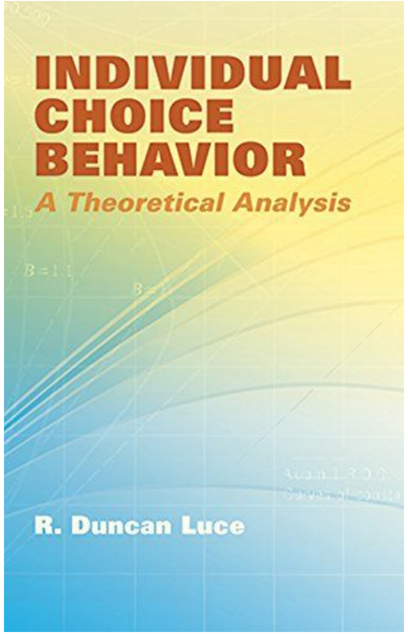

Design Galleries [Marks et al. 1997]

- Seminal work on display representative design options on a two-dimensional screen by low dimensional embedding
- Focused on light placement settings
    - Didn't mention anything on the preference theory :(
- A few follow up works
    - Text editing preview: M Terry et al CHI'02
    - Web design: B Lee et al CHI' 10
    - Crowd-powered visual design Y Koyama et al UIST'14
    - Interface design: J Dudley et al CHI'19
    - Input devices: Y Liao CHI' 21 EA
    - … too many!

Reflections

Cognitive Illusions in Sequential Design Preferences

Cognitive Illusions in Sequential Design Preferences

- *Anchoring*: subjective preferences are dominated by earlier experiences
  - A prompt creates in the subject's mind, at least temporarily, the possibility that the uncertain quantity could be either above or below the prompt
  - Result in psychophysical discrimination errors
  - **User preference function is a conditional distribution:** $g(\mathbf{x}|u)$

Cognitive Illusions in Sequential Design Preferences

- *Anchoring*: subjective preferences are dominated by earlier experiences
  - A prompt creates in the subject's mind, at least temporarily, the possibility that the uncertain quantity could be either above or below the prompt
  - Result in psychophysical discrimination errors
  - **User preference function is a conditional distribution:** $g(\mathbf{x}|u)$
- *Drifting*: subjective preference changes over time in an open-ended design task
  - The mesh optimizer fails to converge to an optimal design because of invalid feedbacks (garbage in, garbage out)
  - **User preference can be time dependent:** $g(\mathbf{x}, t_1|u) \neq g(\mathbf{x}, t_2|u)$

Cognitive Illusions in Sequential Design Preferences

- *Anchoring*: subjective preferences are dominated by earlier experiences
  - A prompt creates in the subject's mind, at least temporarily, the possibility that the uncertain quantity could be either above or below the prompt
  - Result in psychophysical discrimination errors
  - **User preference function is a conditional distribution:** $g(\mathbf{x}|u)$
- *Drifting*: subjective preference changes over time in an open-ended design task
  - The mesh optimizer fails to converge to an optimal design because of invalid feedbacks (garbage in, garbage out)
  - **User preference can be time dependent:** $g(\mathbf{x}, t_1|u) \neq g(\mathbf{x}, t_2|u)$
- *Yearning*: subjects are eager to produce an expected preference or similar
  - Exploration slows the exploitation process down in order to decode the entire design space
  - **Preference space and design space may only have small overlaps:** $g(\mathbf{x}) \in \mathcal{P}, |\mathcal{P} \cap \mathcal{X}| \approx 0$

Sources of Meshing Complexity

1. Well-formulated mathematics that describes continues settings can be inconsistent in discret settings (No-free-lunch: not all properties are preserved, e.g. Laplace-Beltrami)

$$\sum \theta_i = 2\pi \qquad \sum \theta_i < 2\pi$$

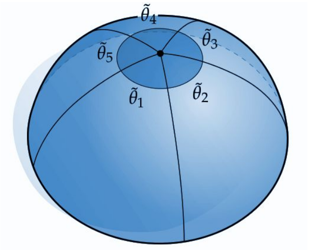

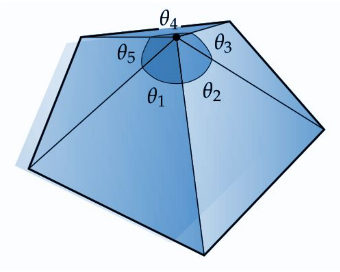

$$(\Delta f)_i = w_i \sum_{ij} w_{ij}(f_j - f_i)$$

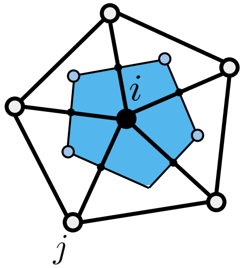

$$w_i = \frac{1}{2A_i}, w_{ij} = \cot\alpha_{ij} + \cot\beta_{ij}$$

$$w_i = \frac{1}{N_i}, w_{ij} = 1$$

$$w_{ij} = \frac{1}{||f_i - f_j||}\left(\tan\left(\frac{\gamma_{ij}}{2}\right) + \tan\left(\frac{\delta_{ij}}{2}\right)\right)$$

$$\mathrm{d}N = \kappa \mathrm{d}f(X)$$

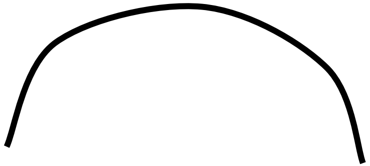

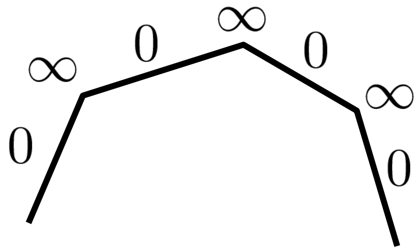

Sources of Meshing Complexity

1. Well-formulated mathematics that describes continues settings can be inconsistent in discret settings (No-free-lunch: not all properties are preserved, e.g. Laplace-Beltrami)

2. Edge cases of geometry processing largely exists and developers failed in handling it well or expensive to give a fix (hard to retrieve the knowledge even for the developer).

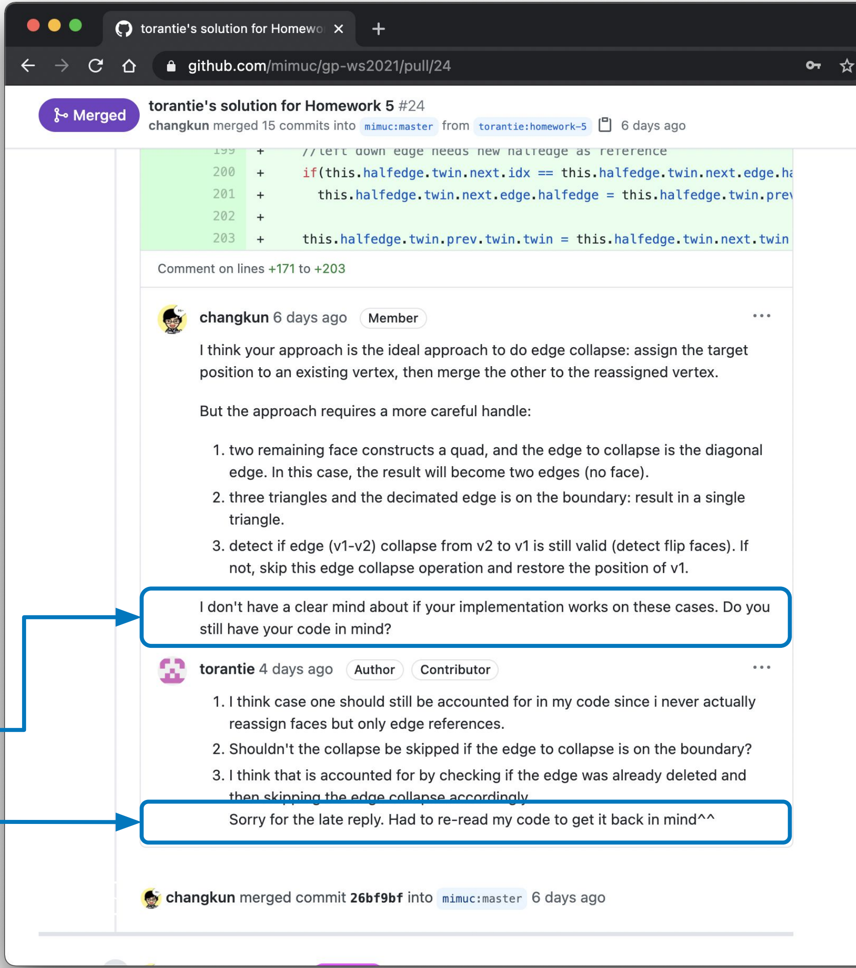

"I don't understand your magic here"

"Neither do I 🤣"

Sources of Meshing Complexity

1. Well-formulated mathematics that describes continues settings can be inconsistent in discret settings (No-free-lunch: not all properties are preserved, e.g. Laplace-Beltrami)

2. Edge cases of geometry processing largely exists and developers failed in handling it well or expensive to give a fix (hard to retrieve the knowledge even for the developer).

3. Ground truth expert opinions (labels) does not exists in modeling (or not widely accepted)

Aside: All ML models are limited by the quality of the ground truth used to train and test

Sources of Meshing Complexity

1. Well-formulated mathematics that describes continues settings can be inconsistent in discret settings (No-free-lunch: not all properties are preserved, e.g. Laplace-Beltrami)

2. Edge cases of geometry processing largely exists and developers failed in handling it well or expensive to give a fix (hard to retrieve the knowledge even for the developer).

3. Ground truth expert opinions (labels) does not exists in modeling (or not widely accepted)

4. Defects exist and slow down the workflow, manual repairing is tedious, but regular users don't care

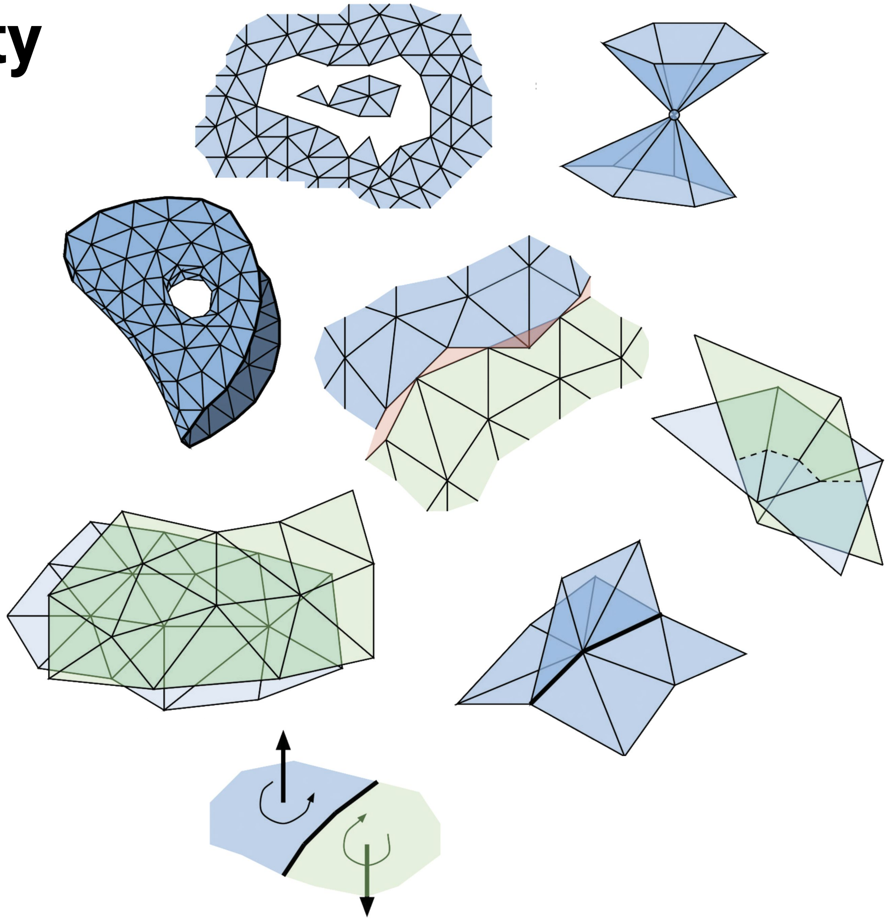

Sources of Meshing Complexity

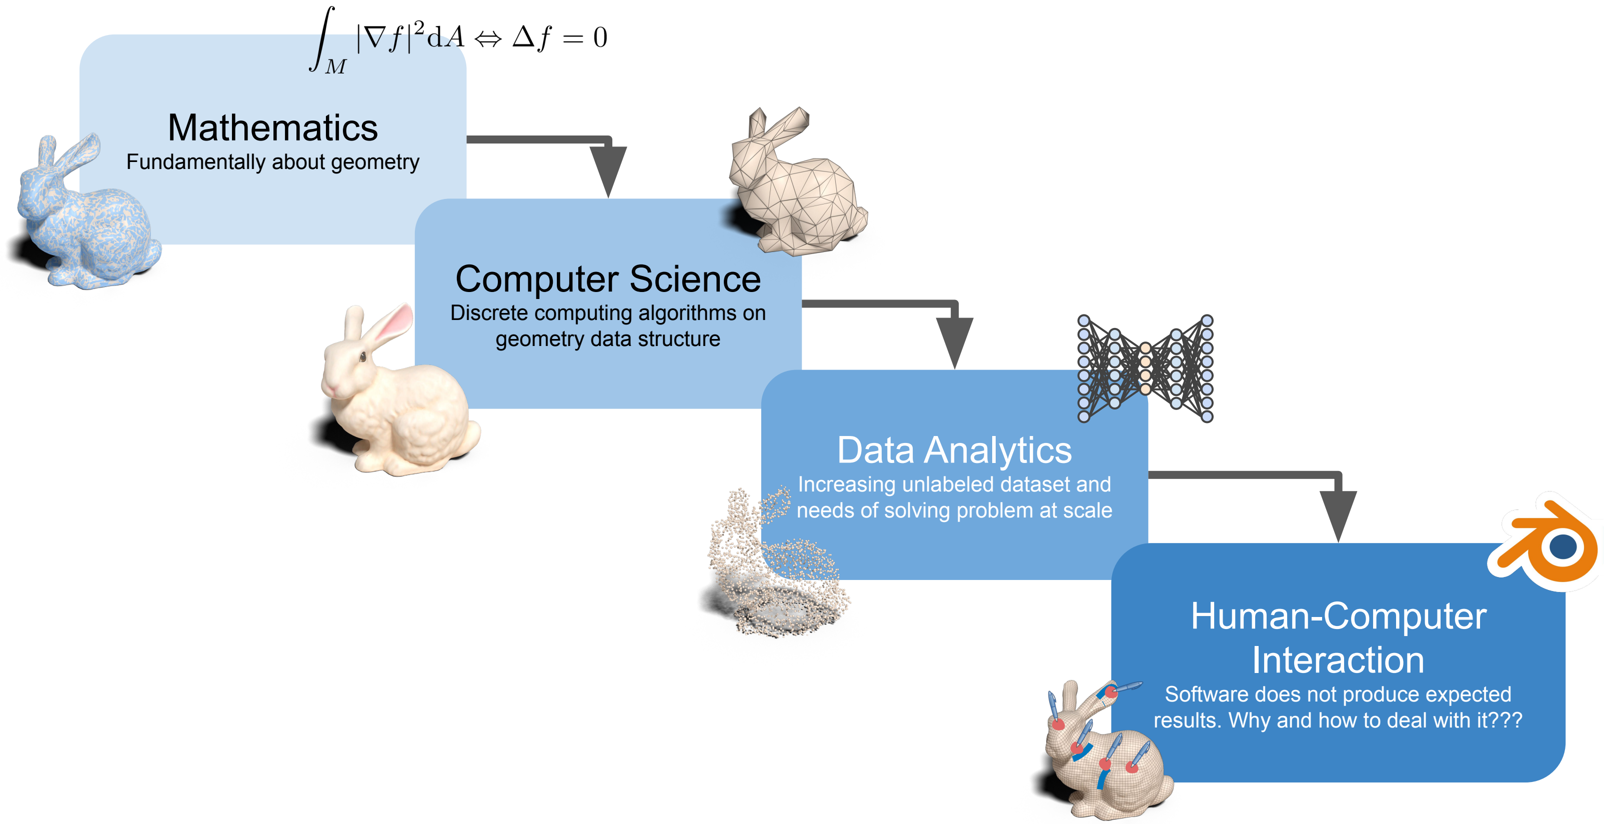

"(So what?) I think theses works are between communities, which community to you want to contribute? (Pick a side)"

-- Prof. Schmidt

When was your last time to do pixel-level adjustment in Photoshop?

How trivial to adjust a pixel?

When was your last time to do vertex-level adjustment in Blender edit mode?

How trivial to adjust a vertex?

Conclusions

- Apparently, all these problems are in HCI (convinced? If not, where? SIGGRAPH? NeurIPS? :-)
  - I believe solving them requires bottom-to-top understanding in geometry and processing
- 3D modeling is a difficult design task in seating in between open-ended and deterministic
  - To the best of my knowledge: all existing HCI theories are not directly applicable here
- The efficiency of 3D modeling workflow does not only depends on algorithms
  - Upstream producer determines characteristics and defects of outputs
  - Downstream consumer determines the requirements on their inputs
- 3D Modeling user interface influences to user's mental model and impacts the experiment
- **Combining subjective data into an objective computation process is a delicate dance**
  - Huge rooms for improvements and still a big challenge

Join the Hacking 👋

- More works on github.com/changkun and github.com/polyred
- Open-sourced and to be open-sourced (ask me if you are interested and one is not open-sourced yet)
    - changkun.de/s/bo: Bayesian optimization in native Go
    - changkun.de/s/ddd: An efficient 3D rasterizer in native Go
    - changkun.de/s/ray: An efficient 3D path tracer in native Go
    - changkun.de/s/win: A cross-platform window management package in Go
    - poly.red/s/geometry: Geometry processing facilities in native Go
    - poly.red/s/geobench: Geometry processing benchmarking facilities in Go
    - poly.red/s/linalg: Algebra facilities for geometry processing in native Go
    - poly.red/s/formats: Wide range formats support in Mesh loader/exporter in Go
    - poly.red/s/raster: Cross-hardware unikernel for rasterization in Go
    - poly.red/s/pbr: Cross-hardware unikernel ray tracing rendering in Go
    - poly.red/s/tree: A 3D engine in Go
    - … and more!

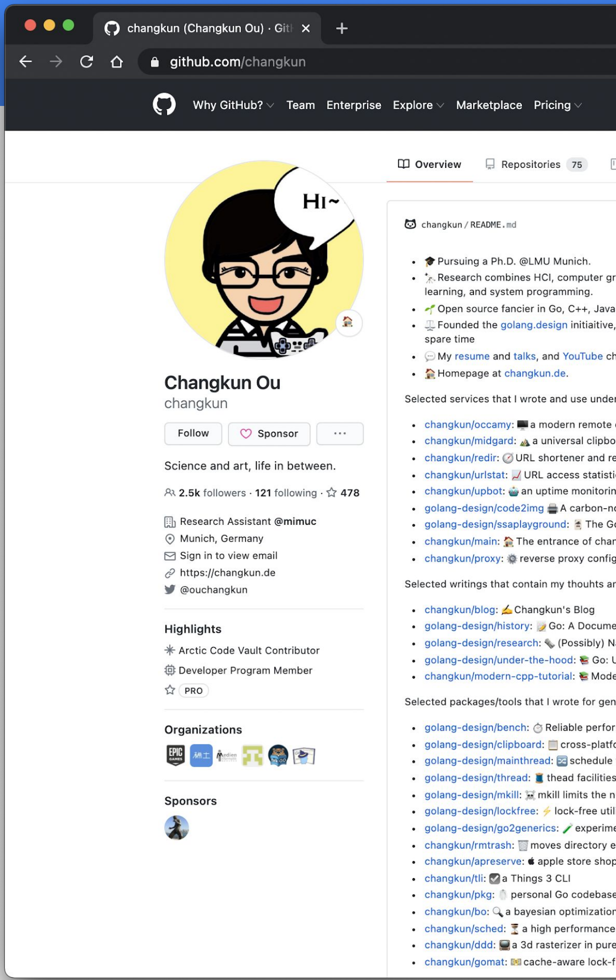

Discussions

- **Better HCI models: (Almost) all existing HCI theories are not applicable**
  - How to better combine users' subjective data into an objective algorithm computation process?
  - Are there any more HCI models that describes and learn human preferences?
- **Efficient preference query:** User should always take the lead but ask feedback is (super) inefficient
  - Can user preference be encoded in an open-ended design task?
  - Search a design as quick as possible (machine side)
- **Time-varying preference learning:** The assumption of users' preference does not change is not valid
  - What can we do about it?
- **Crowdsourcing:** End consumers may not care about high quality mesh
  - How to involve their opinion in real-time? GWAP? [Ou et al 2019]
- **Better coding experience in research**
  - Huge code workload for a simple idea, what can we improve about this?
  - Open-source helps gather crowd wisdom, but I keep doing research as a single coder on a project (almost impossible to communicate low-level details, it is not just calling APIs). How can we improve about this?

Delicate Dance:

Preferences in Interactive Meshing

Changkun Ou
changkun.de/s/polyred5star

@changkun

IDC 2021 Spring
Virtual Event
Munich
Mar 15, 2021