Intelligent Tools for Creative Graphics
by Crowdsourcing

Yuki Koyama

National Institute of Industrial Science and Technology (AIST), Japan

Yuki Koyama

https://koyama.xyz/

Yuki Koyama is a researcher at **National Institute of Advanced Industrial Science and Technology (AIST)**, where he is a member of Media Interaction Group. He received his Ph.D. from The University of Tokyo in 2017, advised by Takeo Igarashi. His research fields are mainly **computer graphics** and **human-computer interaction**. In particular, he is interested in enhancing various design activities by employing computational techniques such as mathematical optimization.

Agenda (25min)

- **Introduction**: How Tools Can Be Intelligent by Crowdsourcing
- **Basics**: Crowdsourcing and Related Concepts
- **Formulation**: Perceptual Feedback from Crowds
- **Intelligent Tools Case 1**: Intelligent Sliders and Suggestions
- **Intelligent Tools Case 2**: Intelligent Automatic Solver
- **Discussions**: Other Types of Intelligence

Introduction

How Tools Can Be Intelligent by Crowdsourcing

Parameter Tweaking is a Common Task

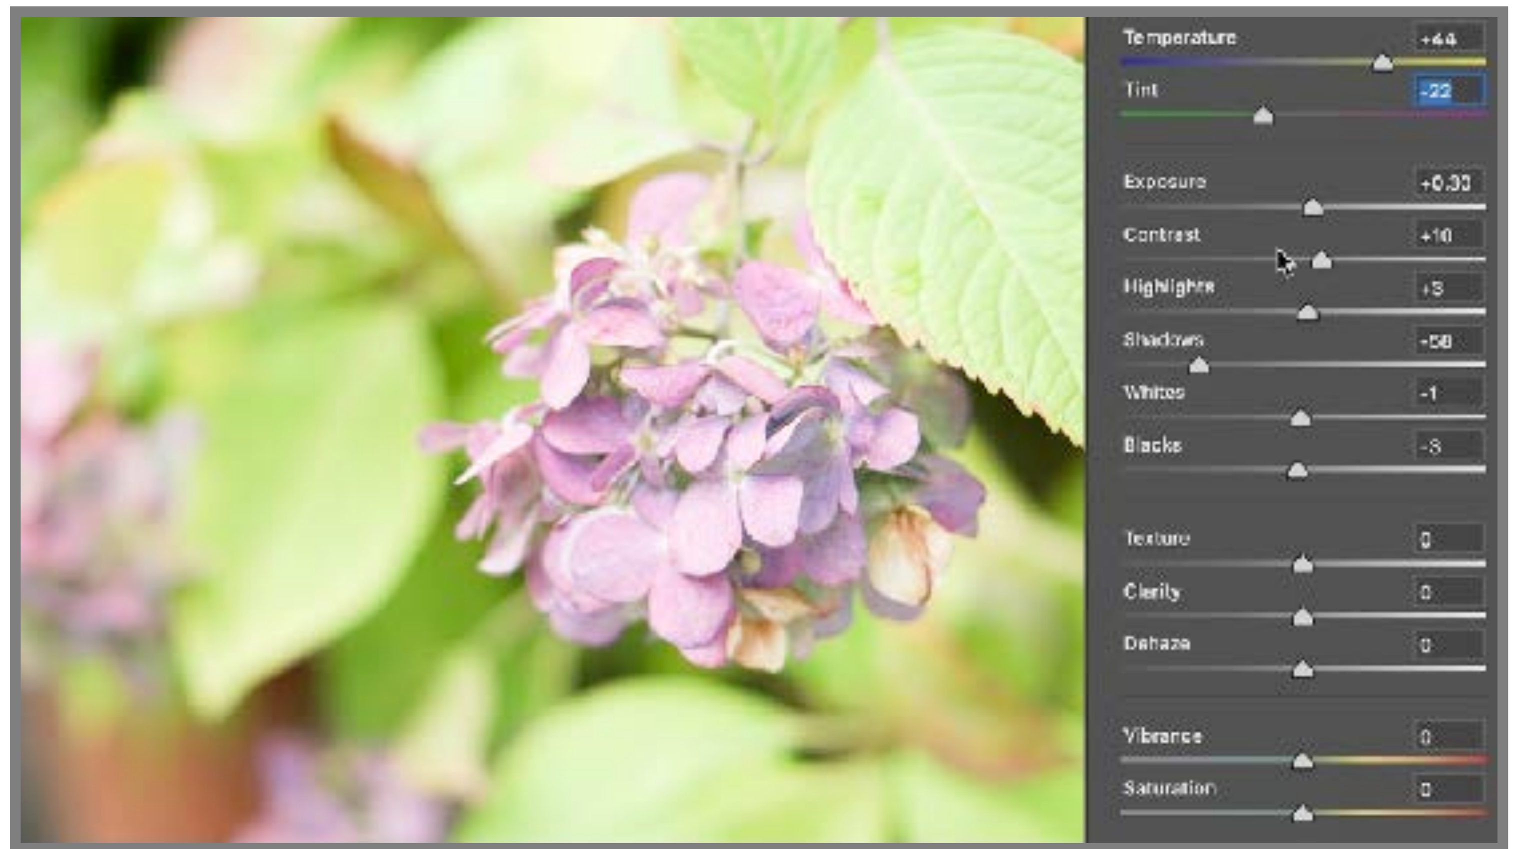

Photo color enhancement

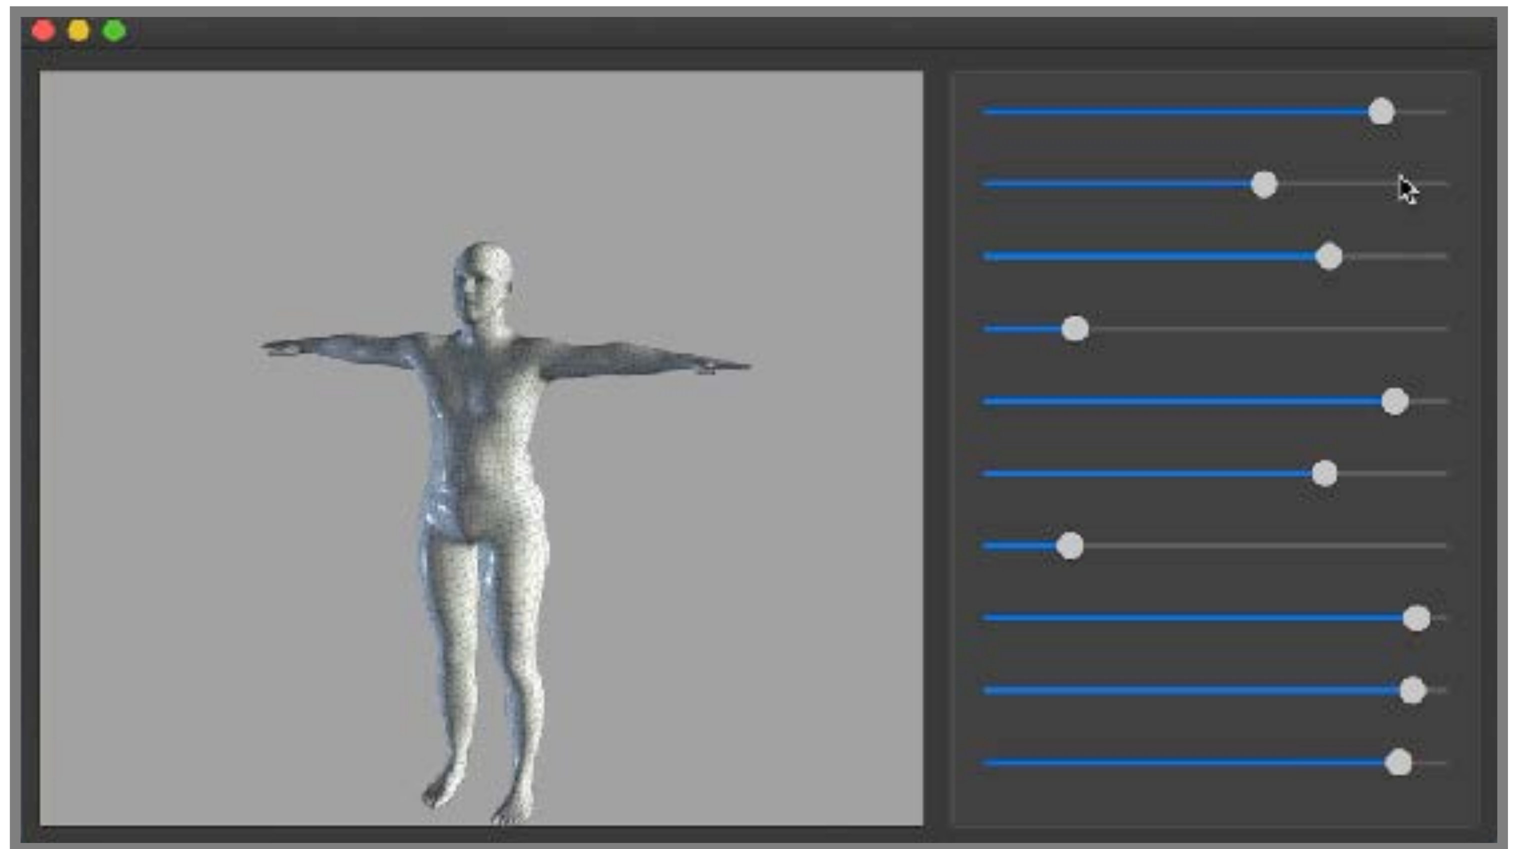

Generative modeling

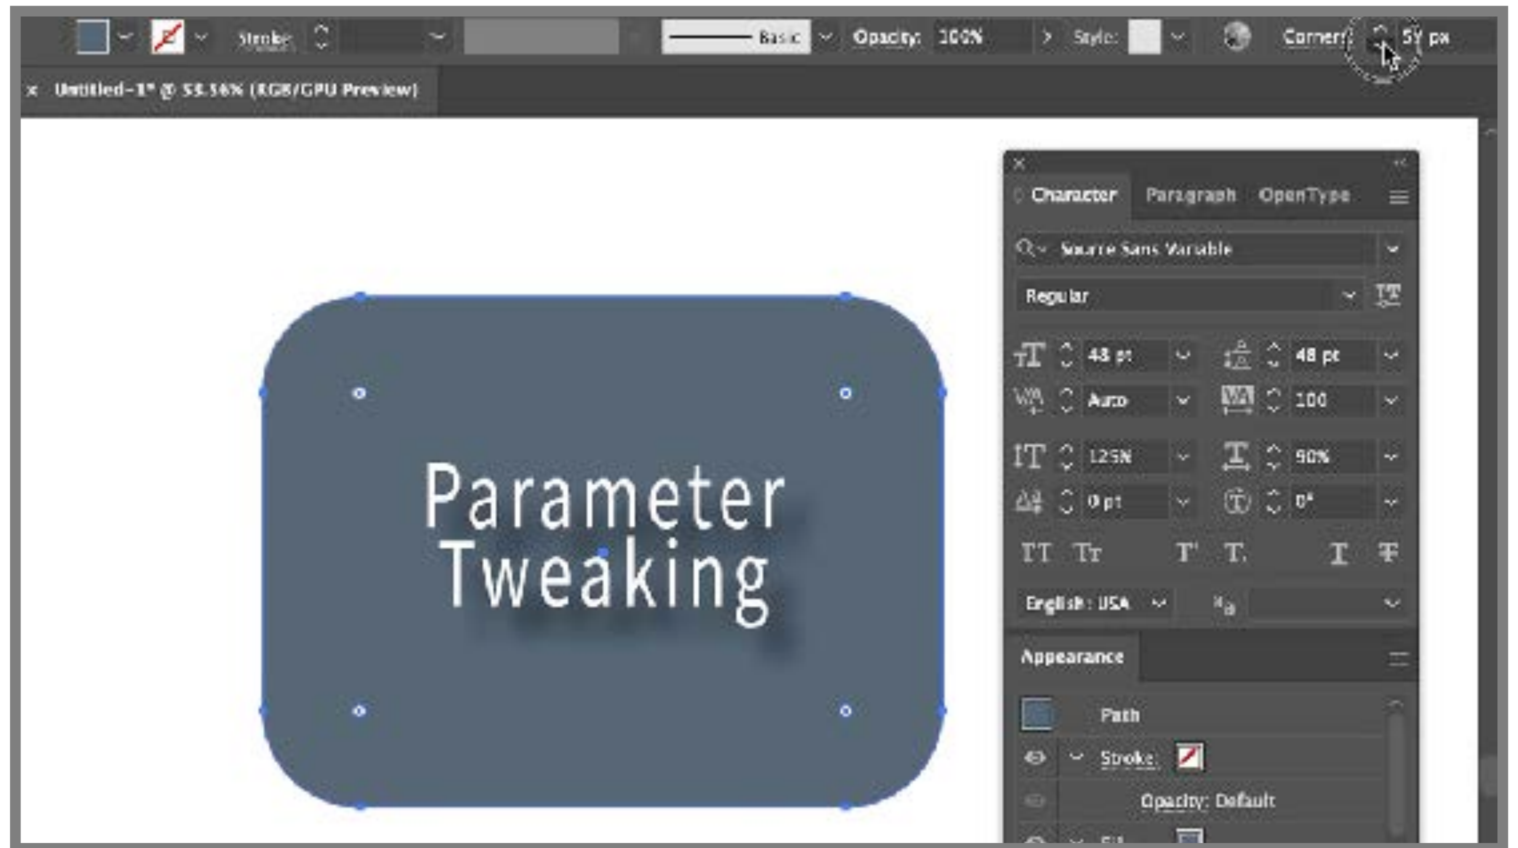

Graphic design

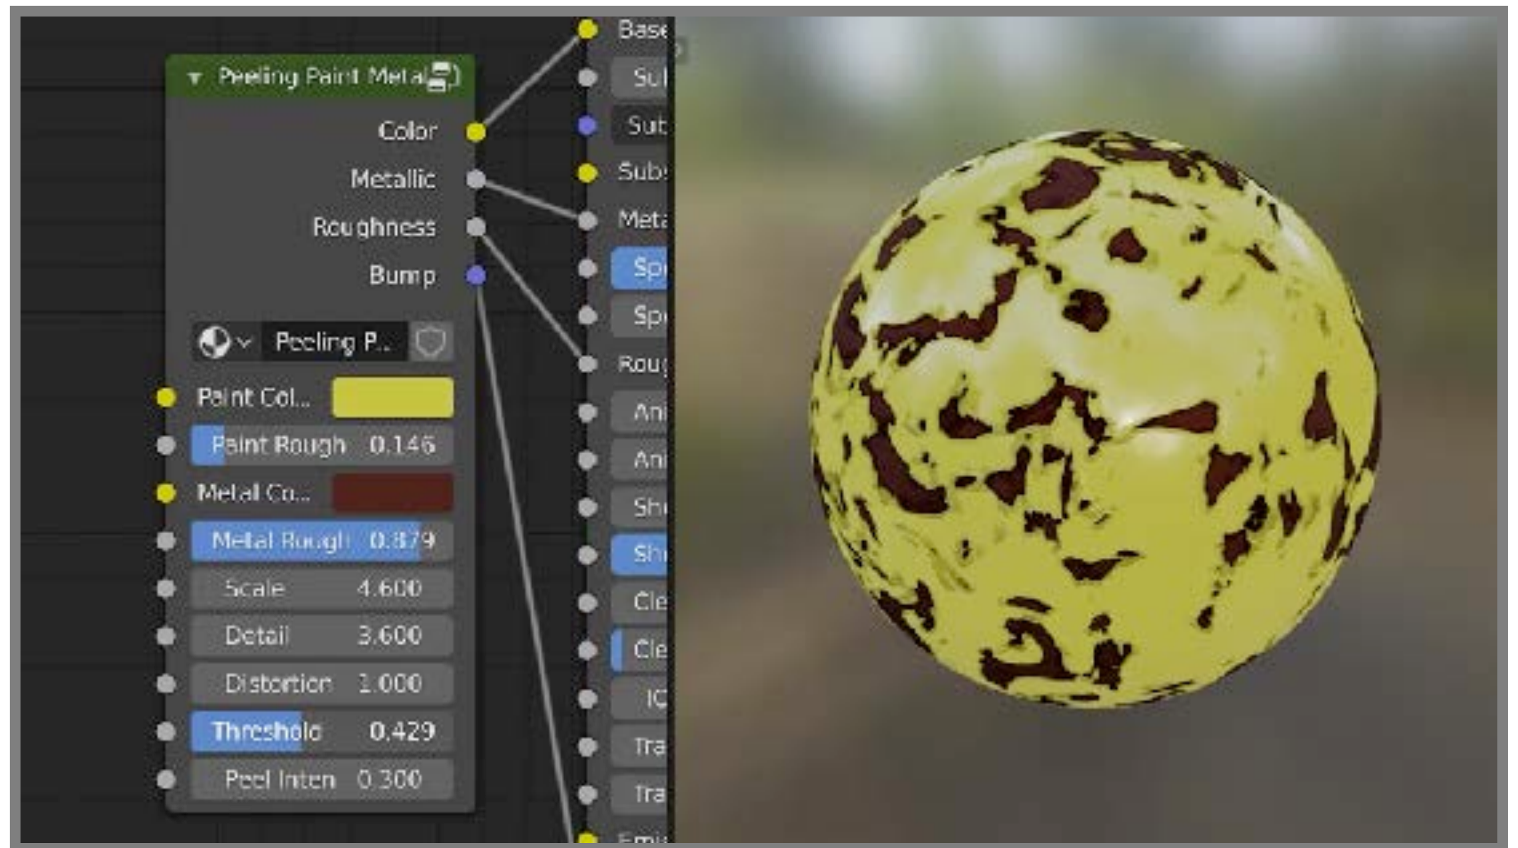

Procedural design

Parameter Tweaking is a Common Task

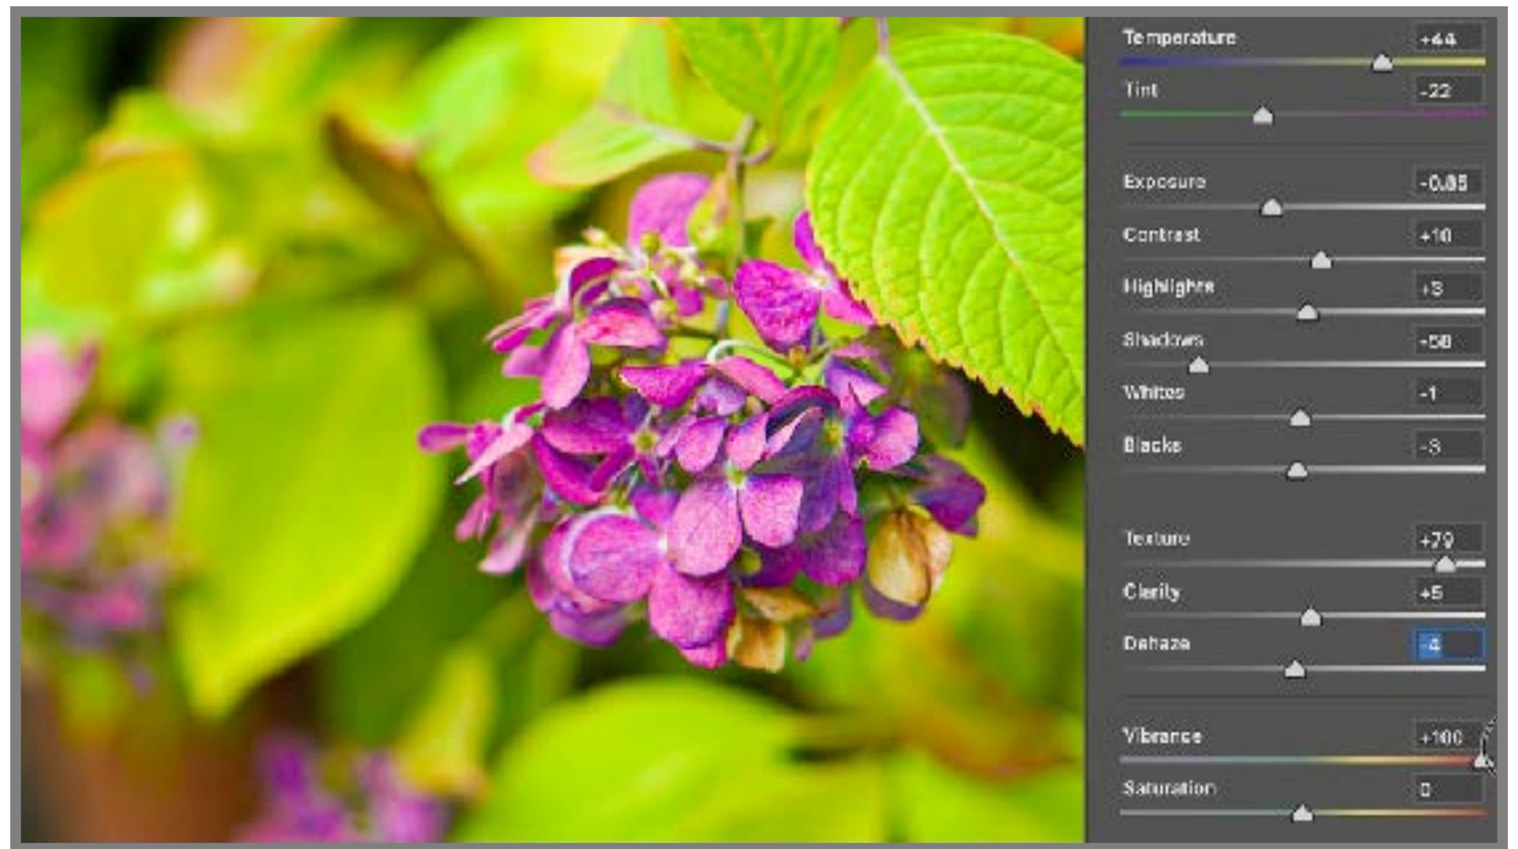

Photo color enhancement

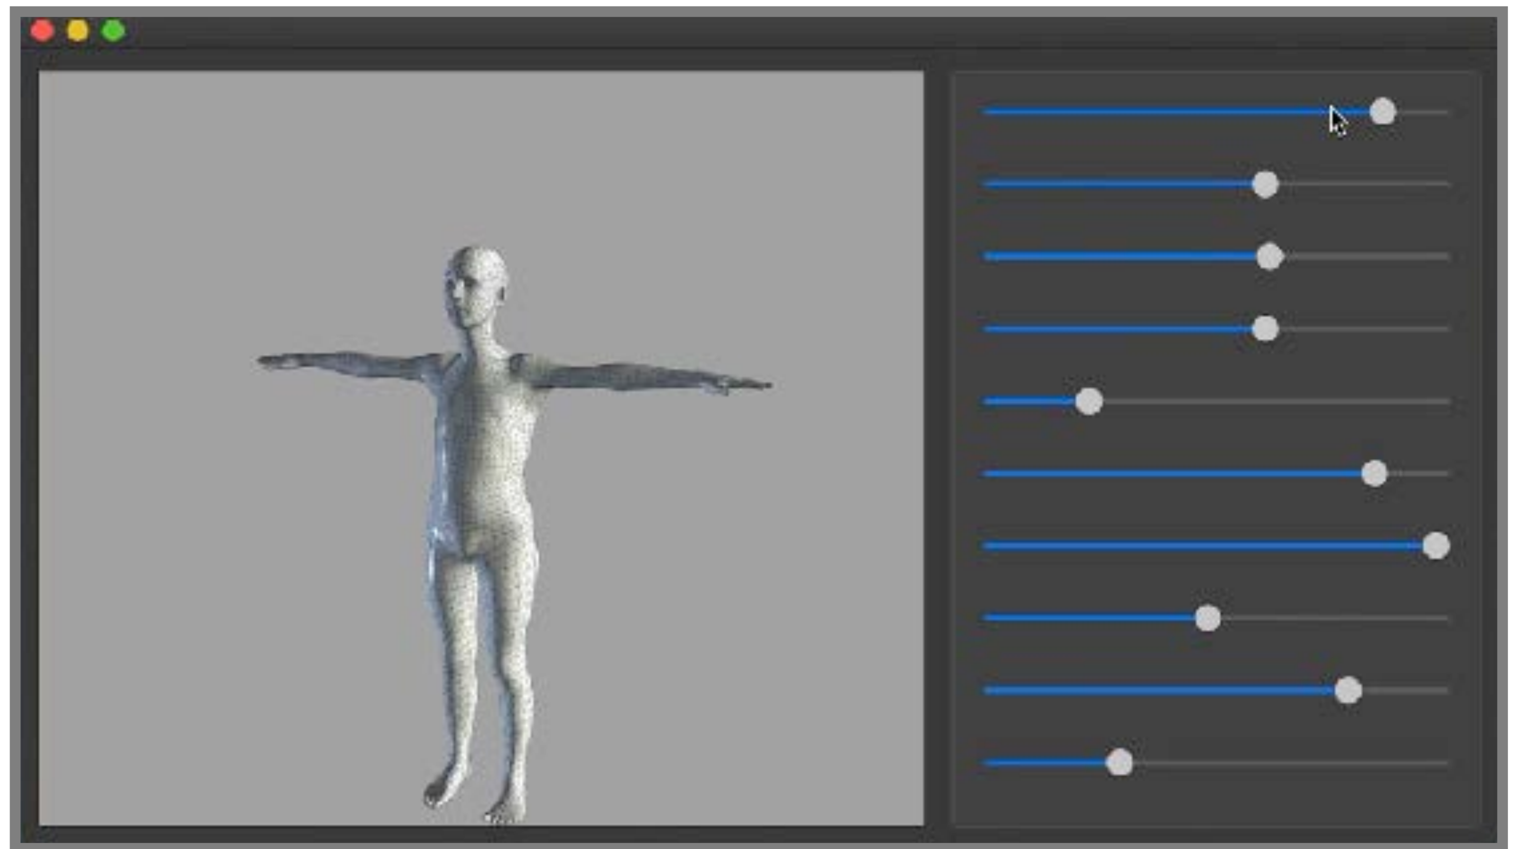

Generative modeling

Graphic design

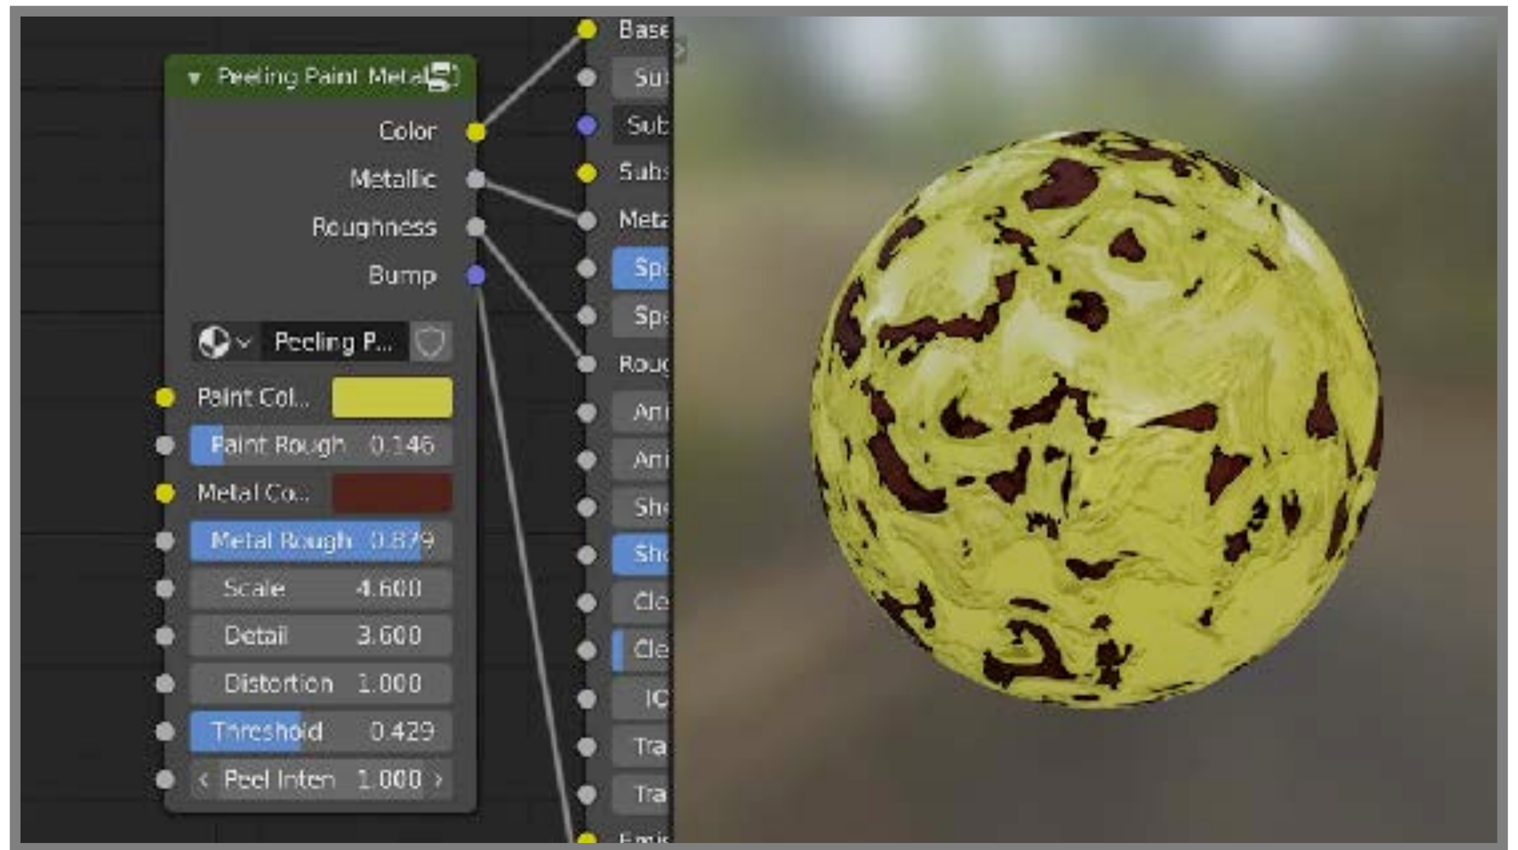

Procedural design

Task: Find the most **preferable** design from the high-dimensional parameter space

Parameter Tweaking with Typical Sliders

Problem:

- This requires many trials and errors, and so this can be **tedious and time-consuming**

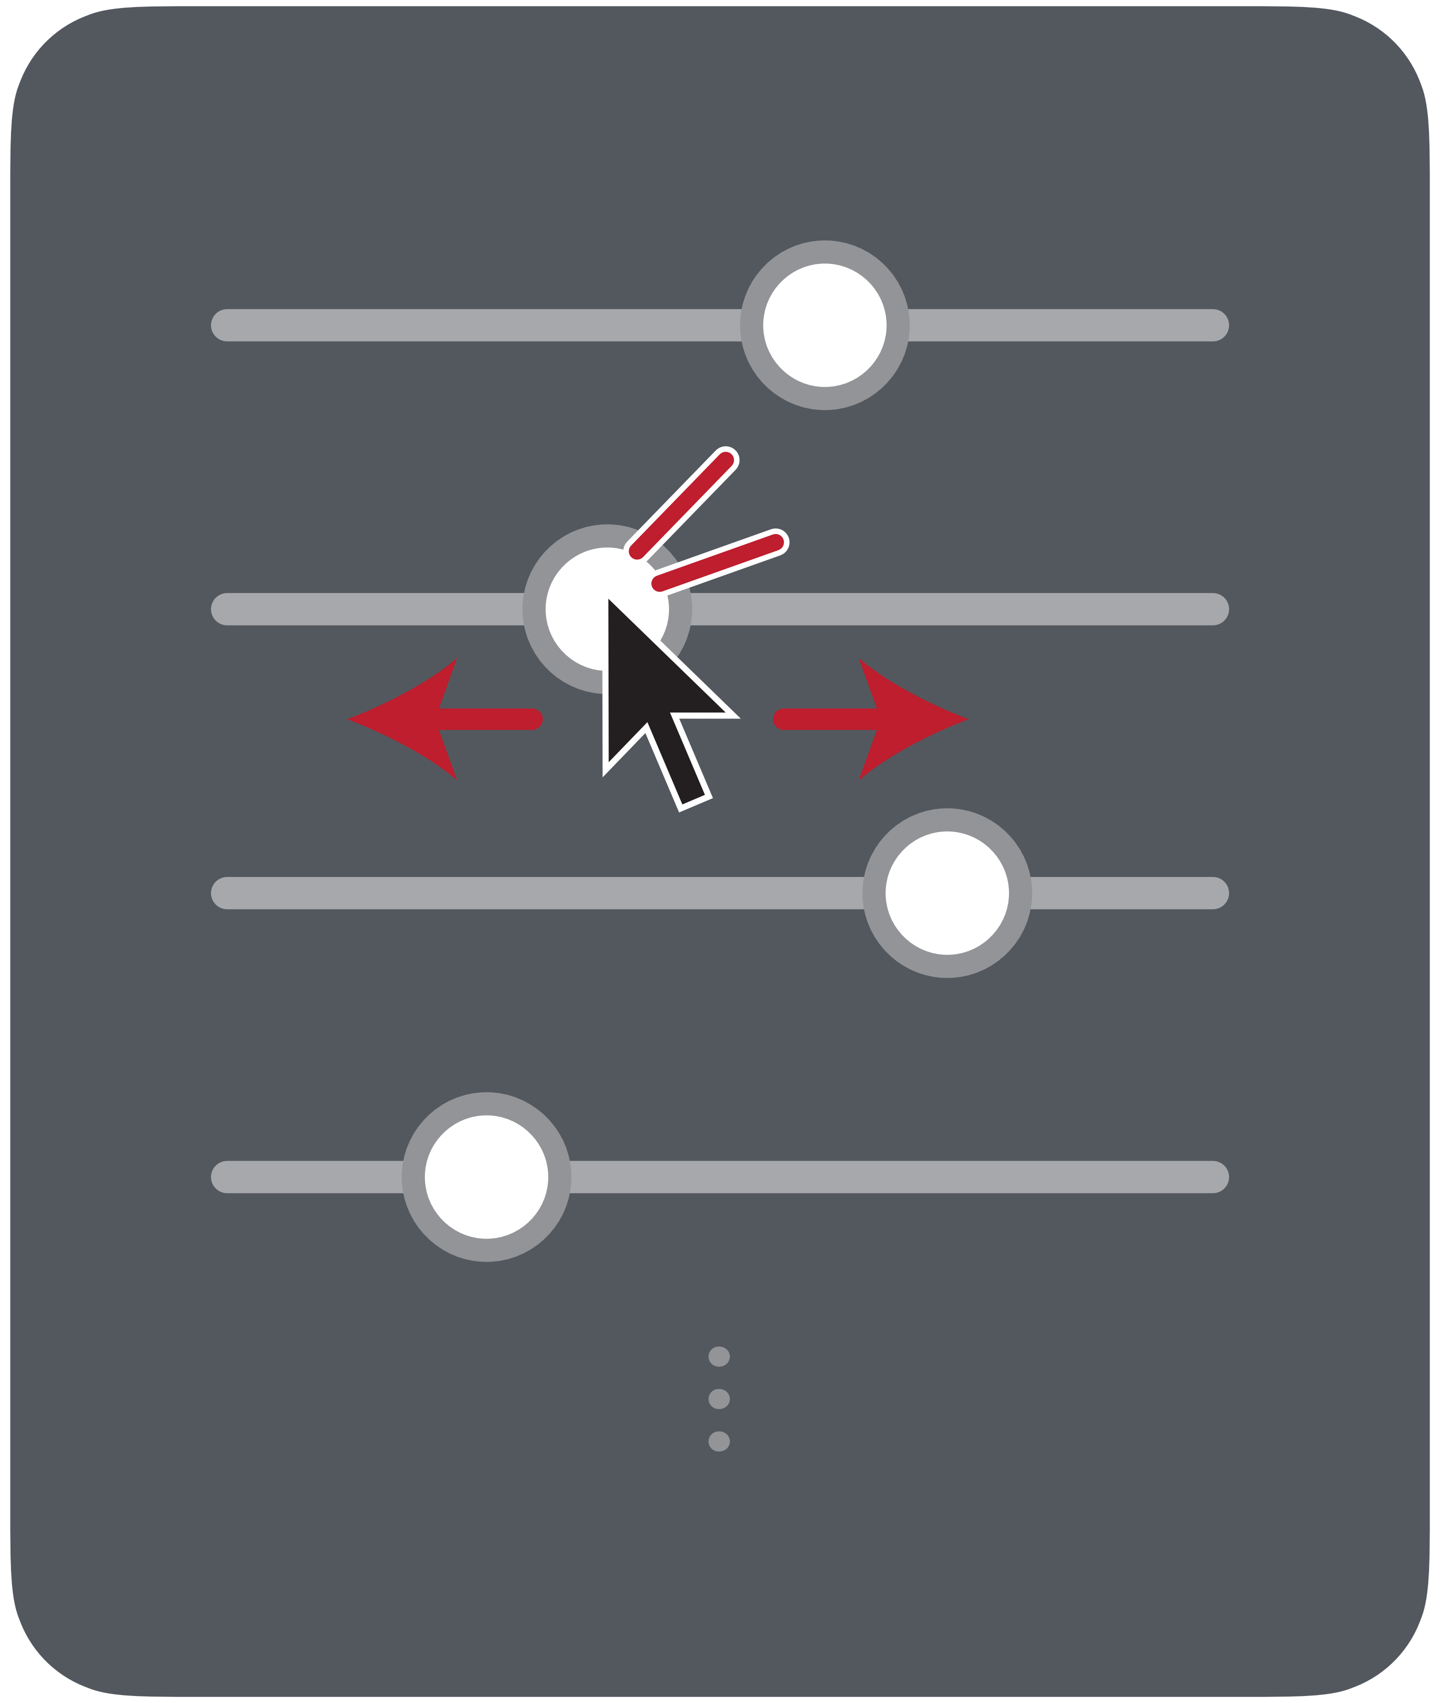

Parameter Tweaking with Typical Sliders

Problem:

- This requires many trials and errors, and so this can be **tedious and time-consuming**

Why not automating via simple scripting?

- **Preference** is based on **human perception**, and so automation is not trivial without some intelligence

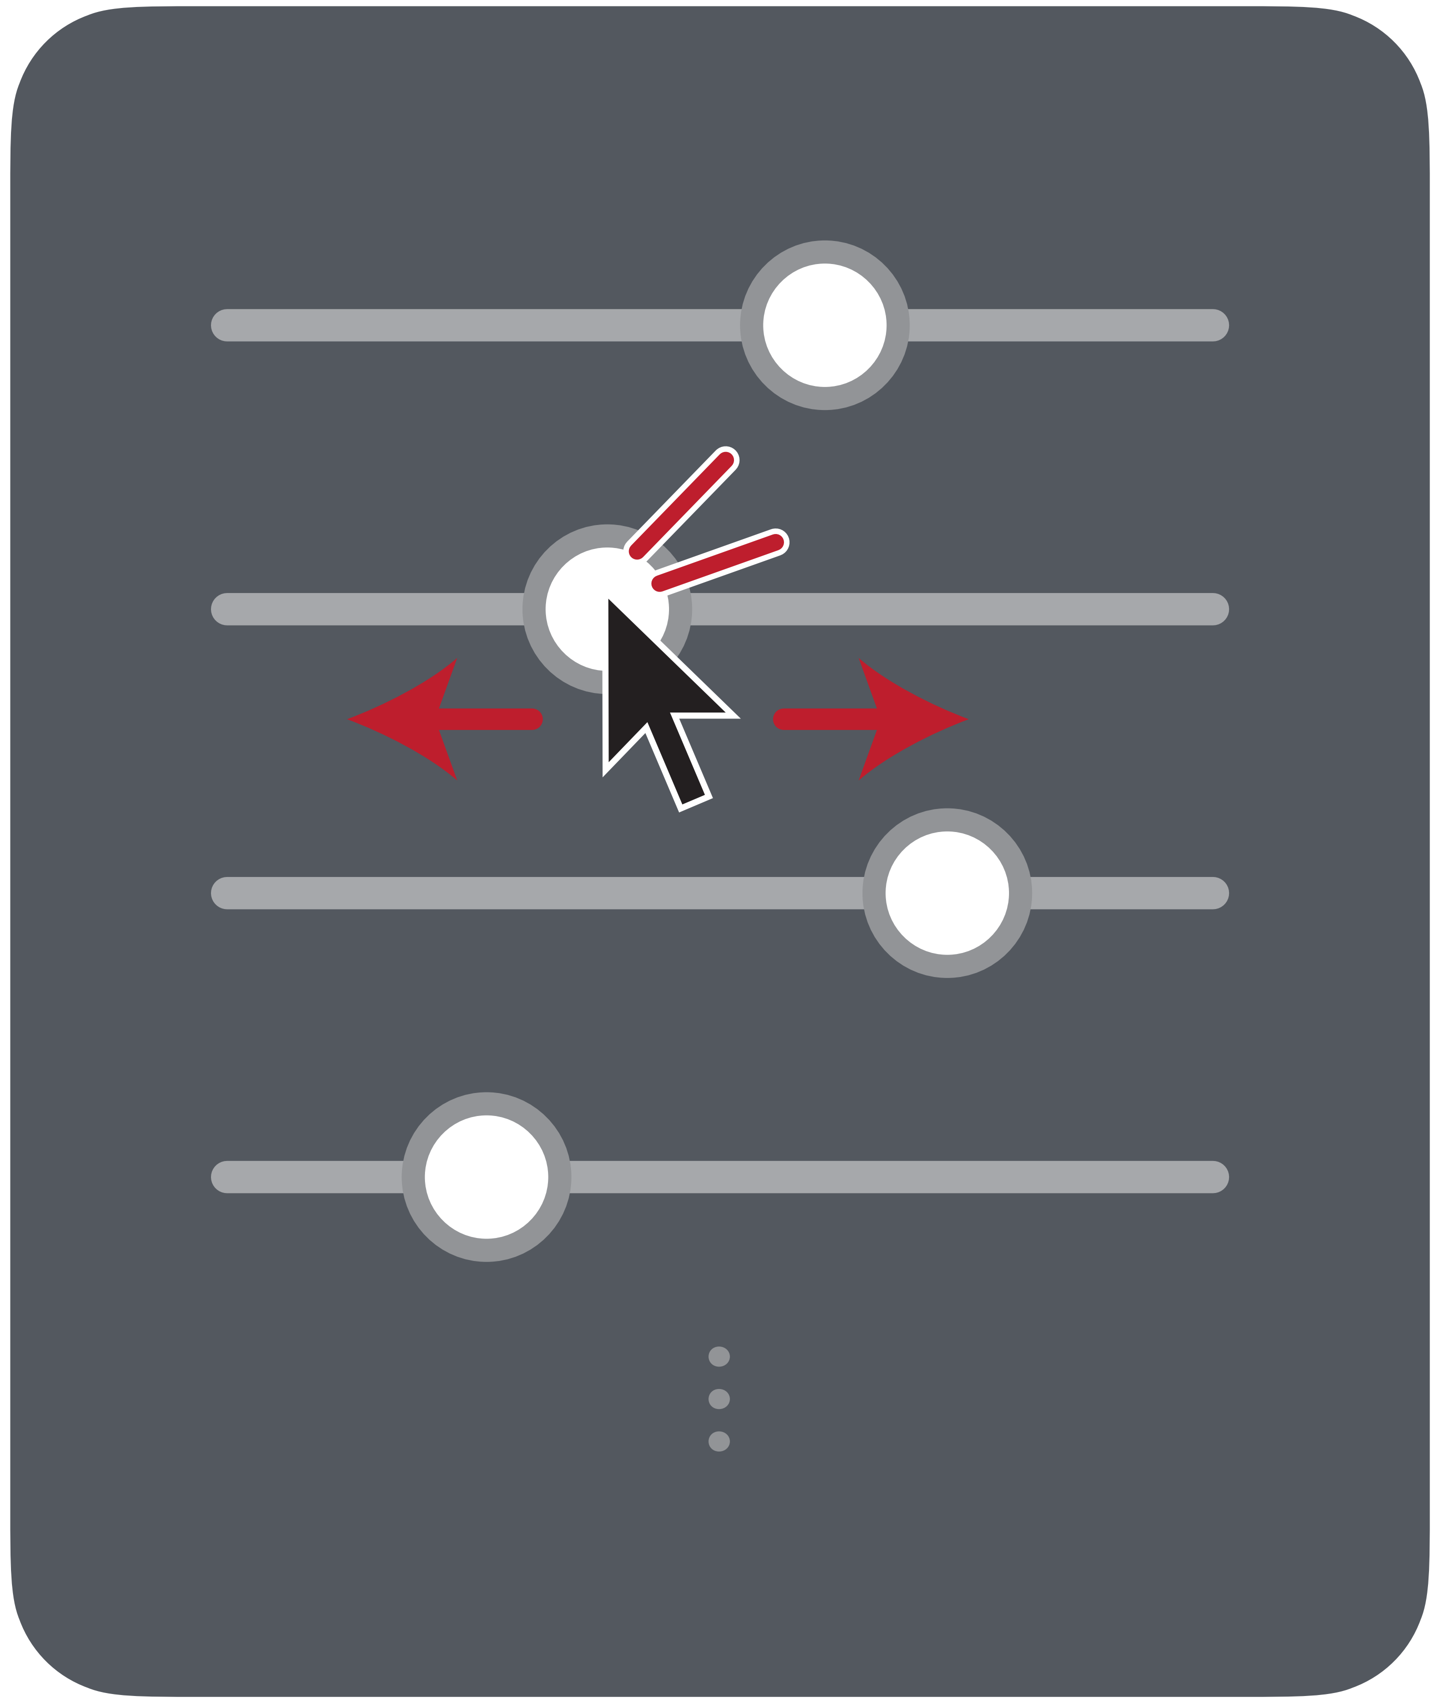

Parameter Tweaking with Typical Sliders

Problem:

- This requires many trials and errors, and so this can be **tedious and time-consuming**

Why not automating via simple scripting?

- **Preference** is based on **human perception**, and so automation is not trivial without some intelligence

Question:

- How can we make tools **intelligent** so that they can handle human perception?

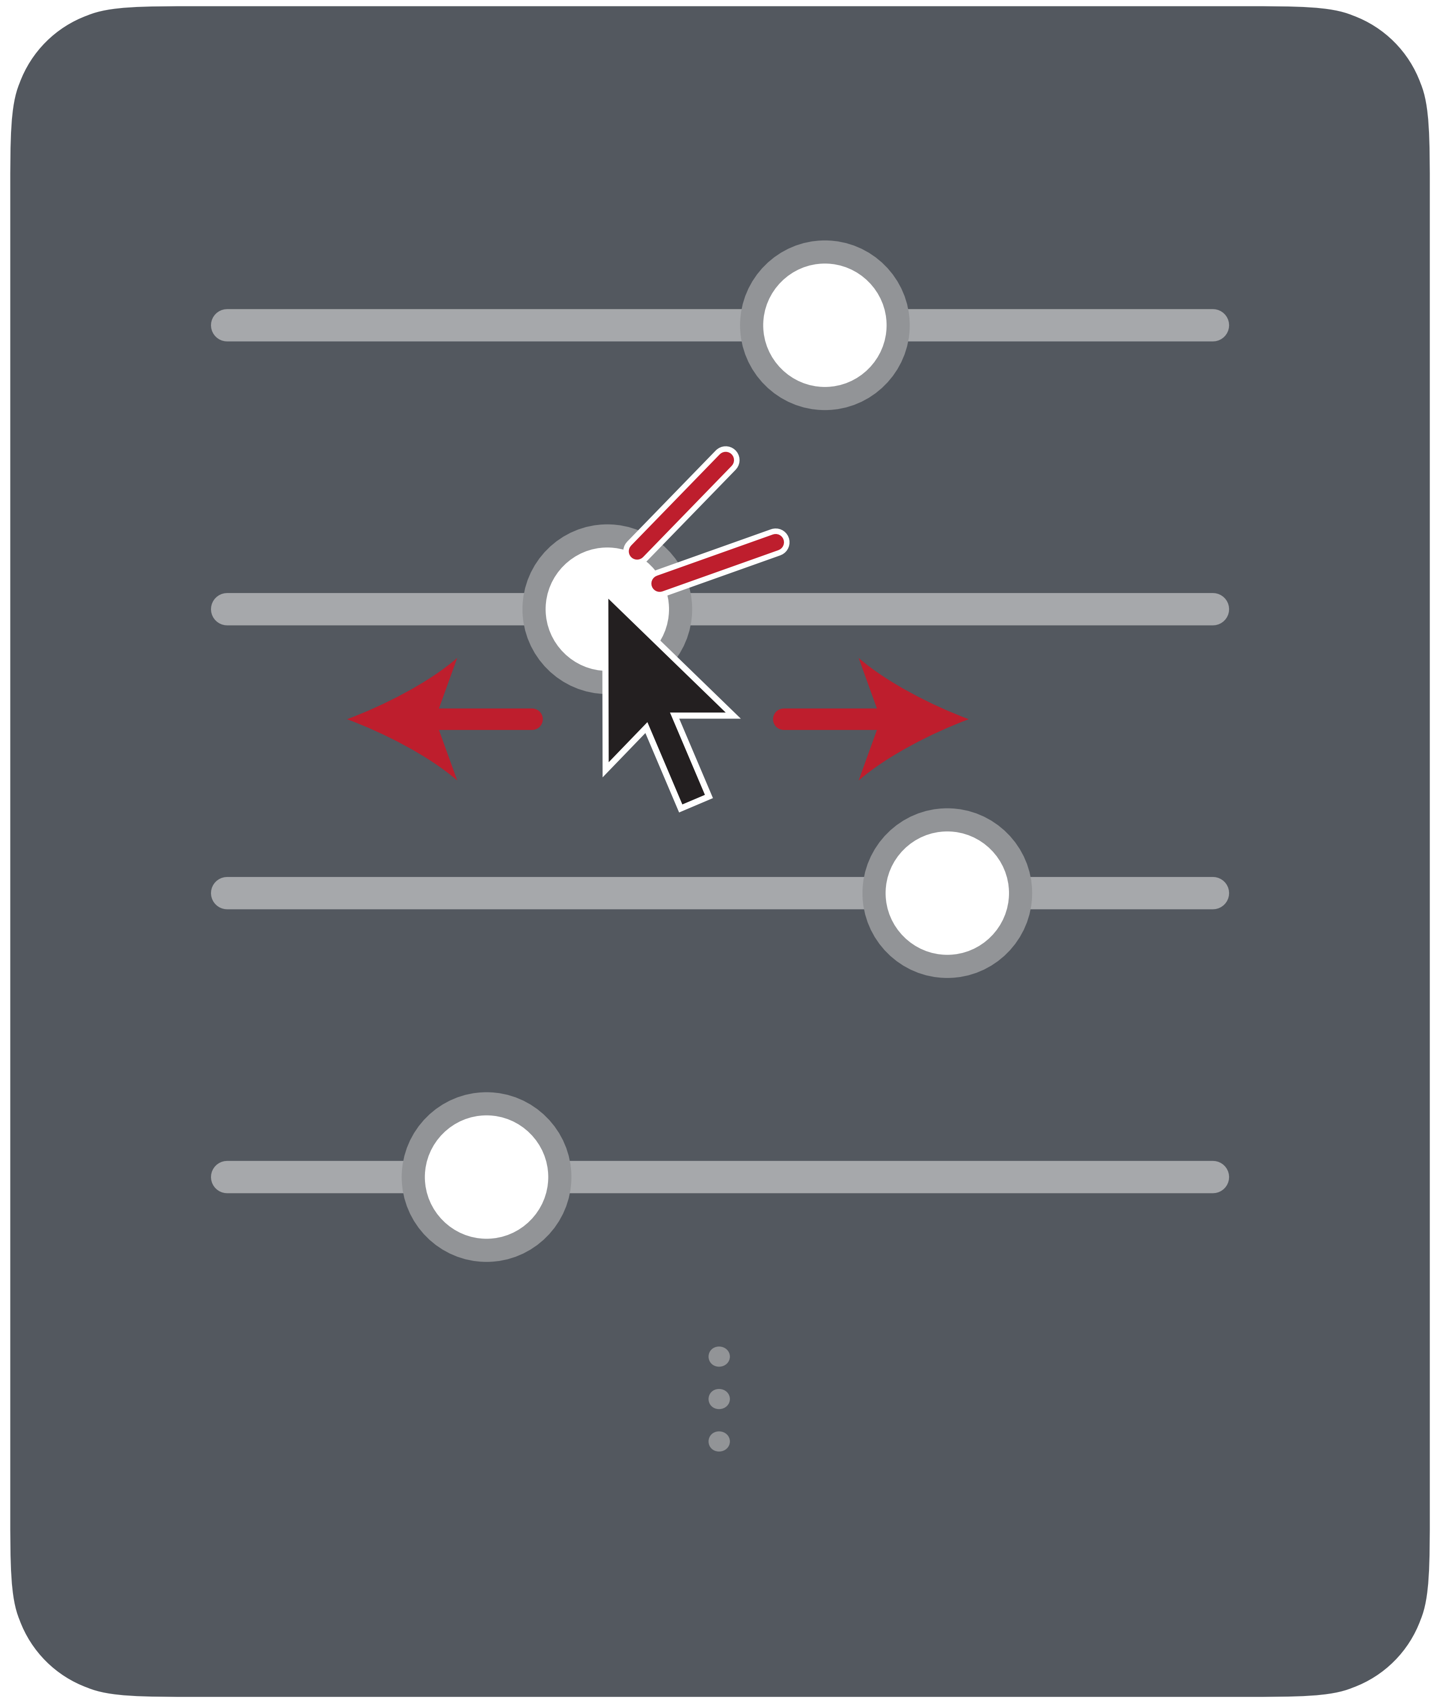

A Solution: Crowdsourcing

Crowdsourcing can be **a software component** of an intelligent tool

- The tool can systematically query crowd workers when necessary

We can incorporate **human intelligence** into the tool

- To quantify perception-related concepts (e.g., preference)

➡ **"Crowd-powered" intelligent tools**

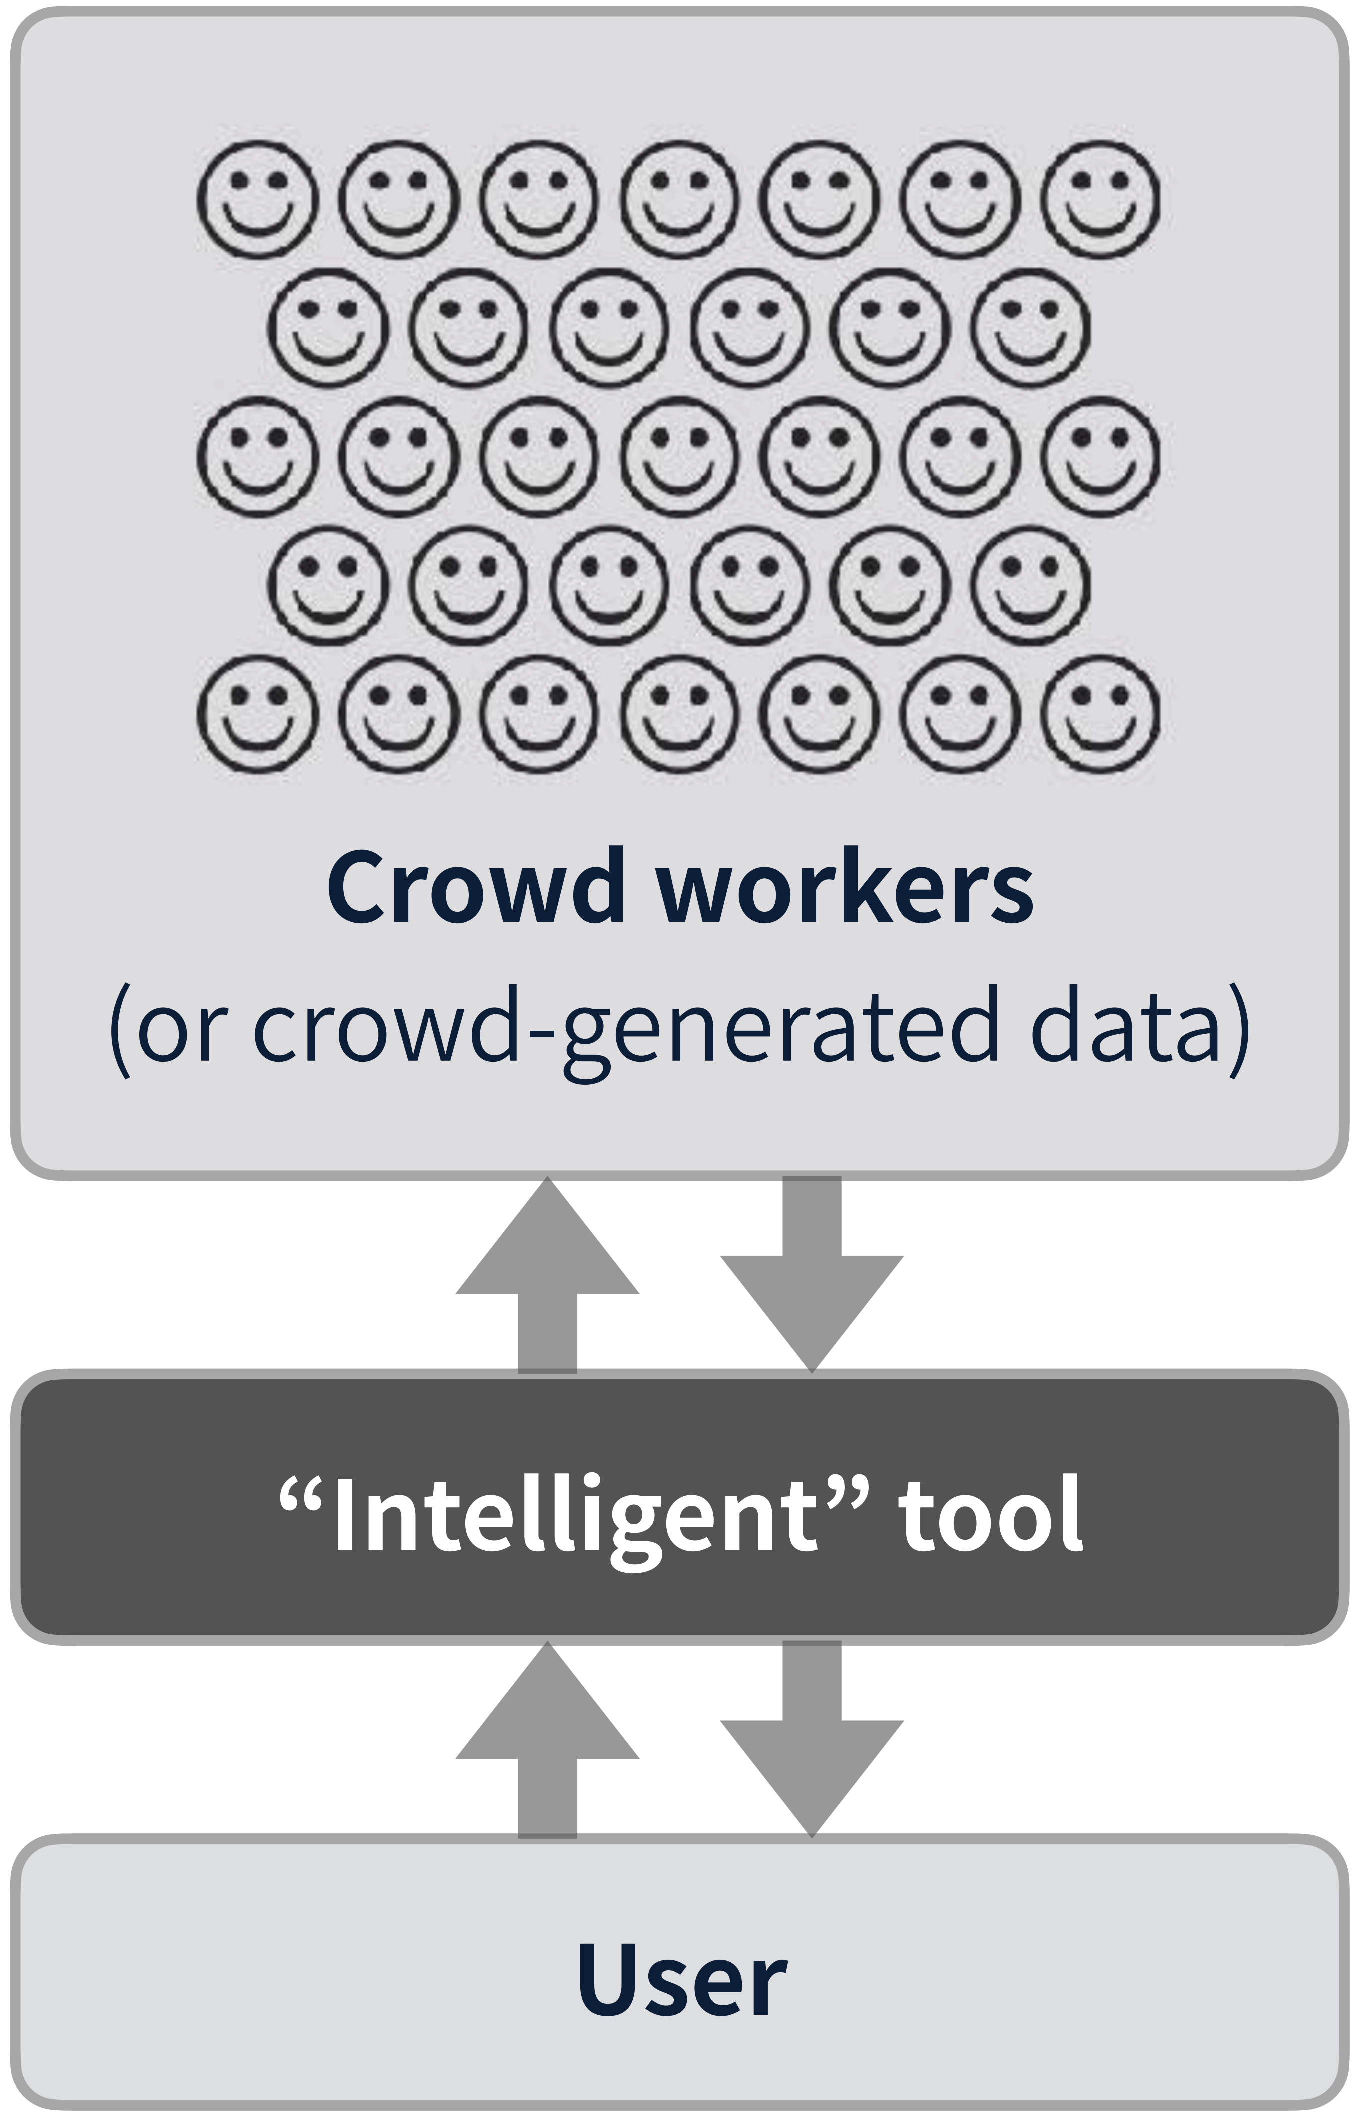

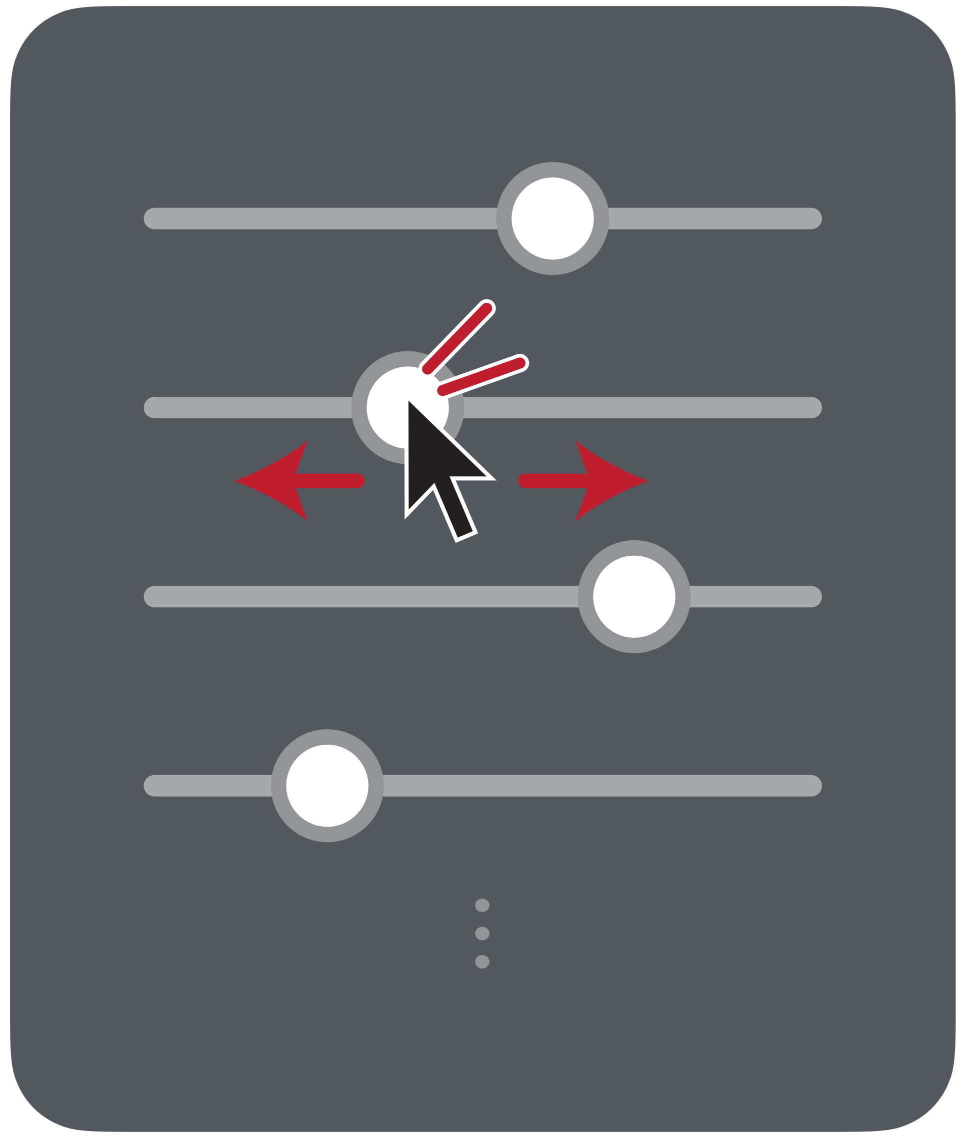

Typical sliders

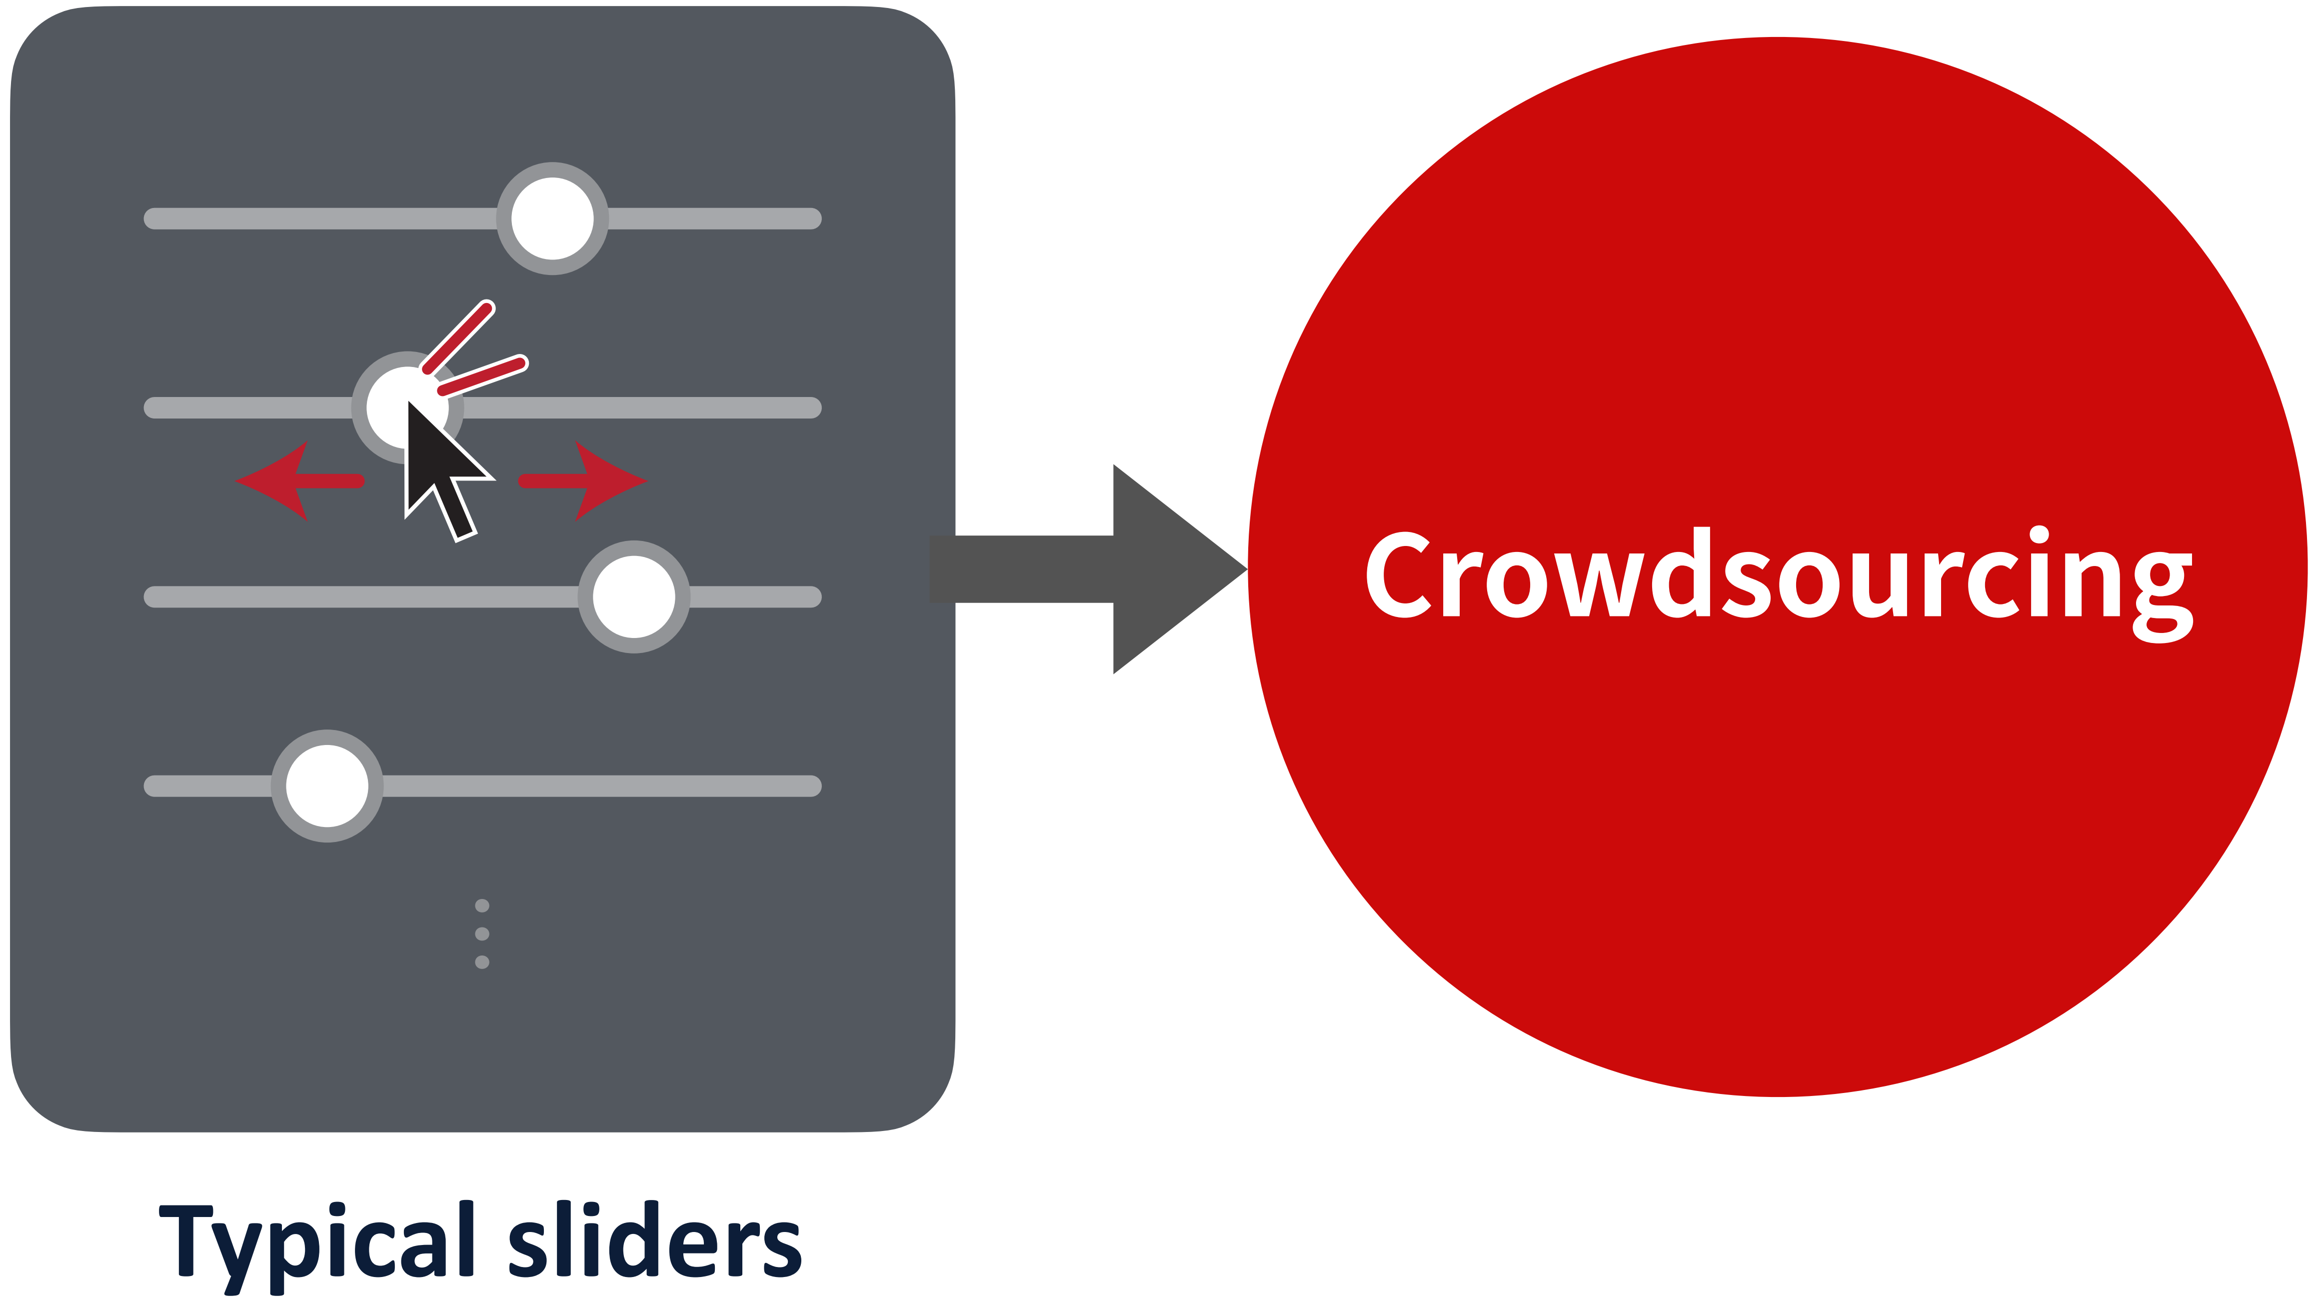
Crowdsourcing
Typical sliders

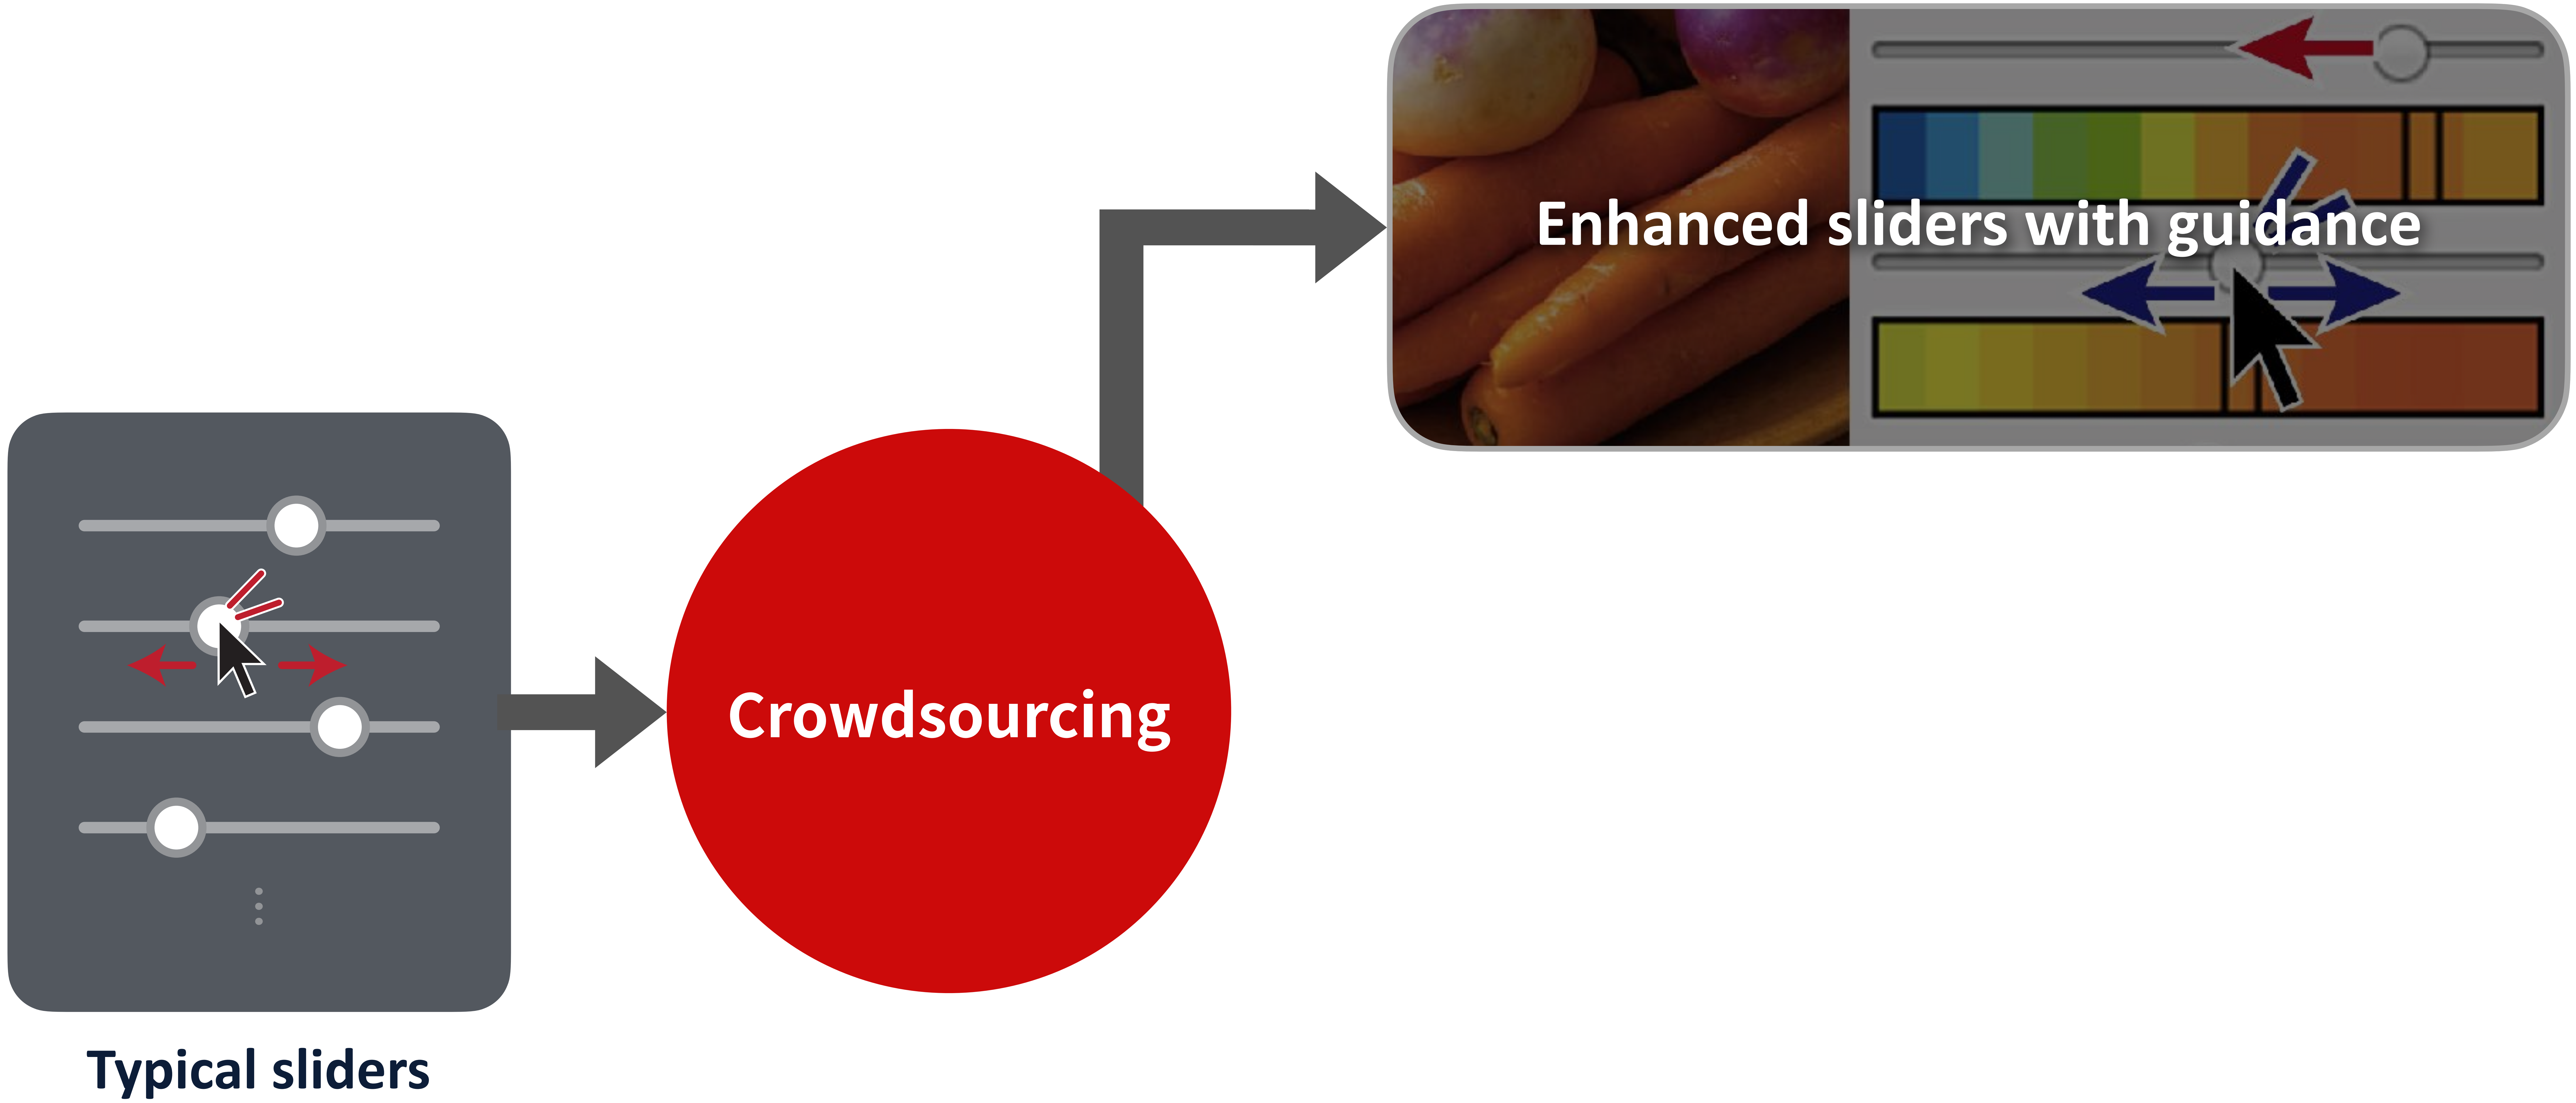
Enhanced sliders with guidance
Crowdsourcing
Typical sliders

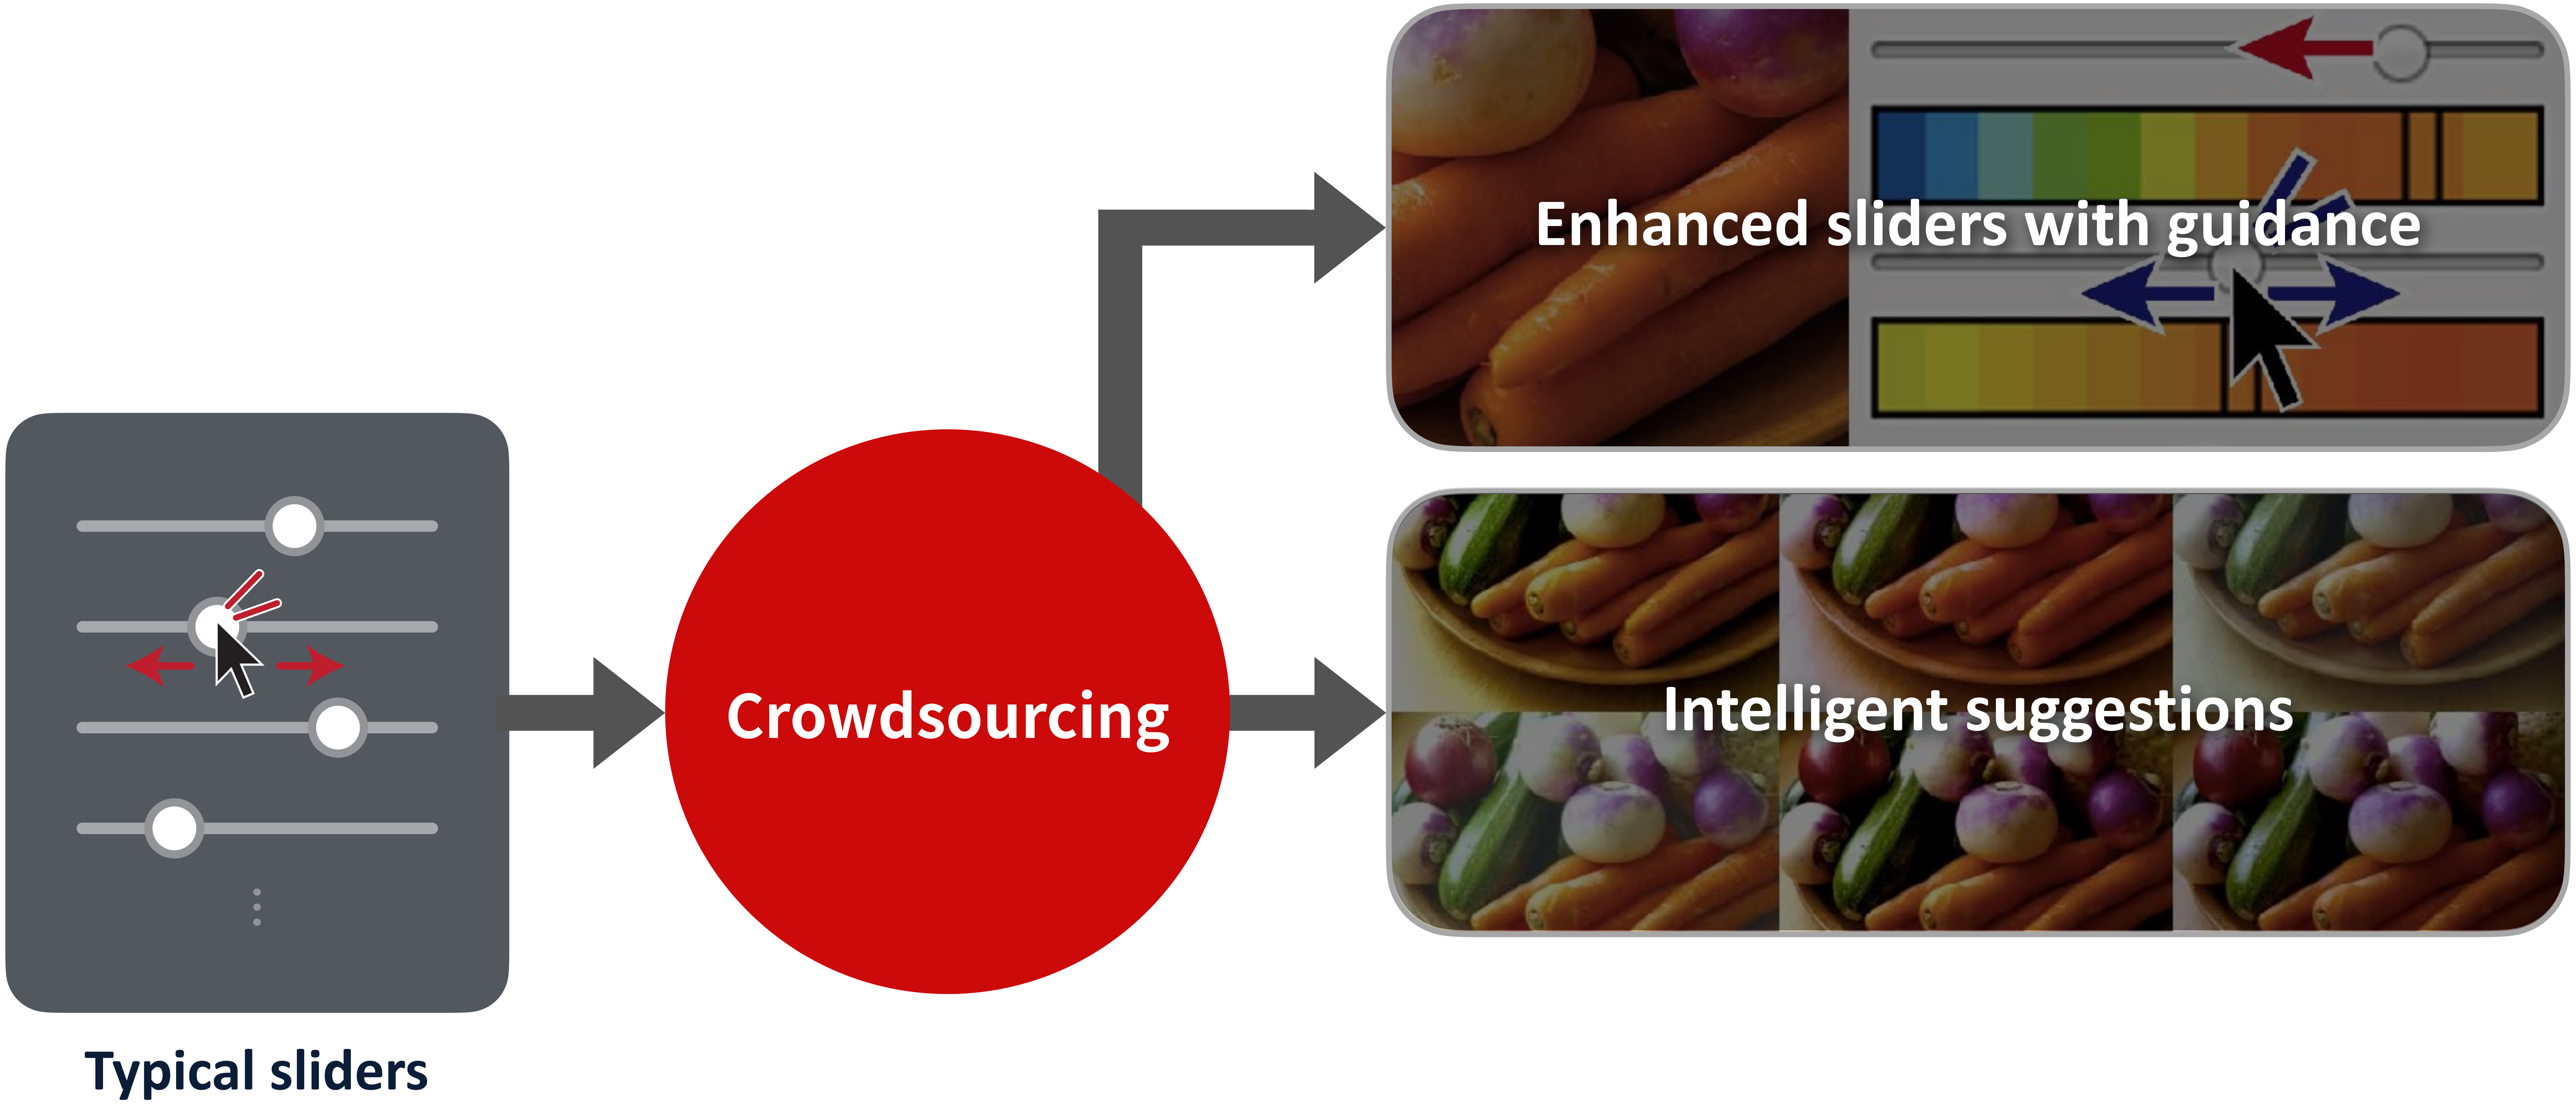
Enhanced sliders with guidance
Crowdsourcing
Intelligent suggestions
Typical sliders

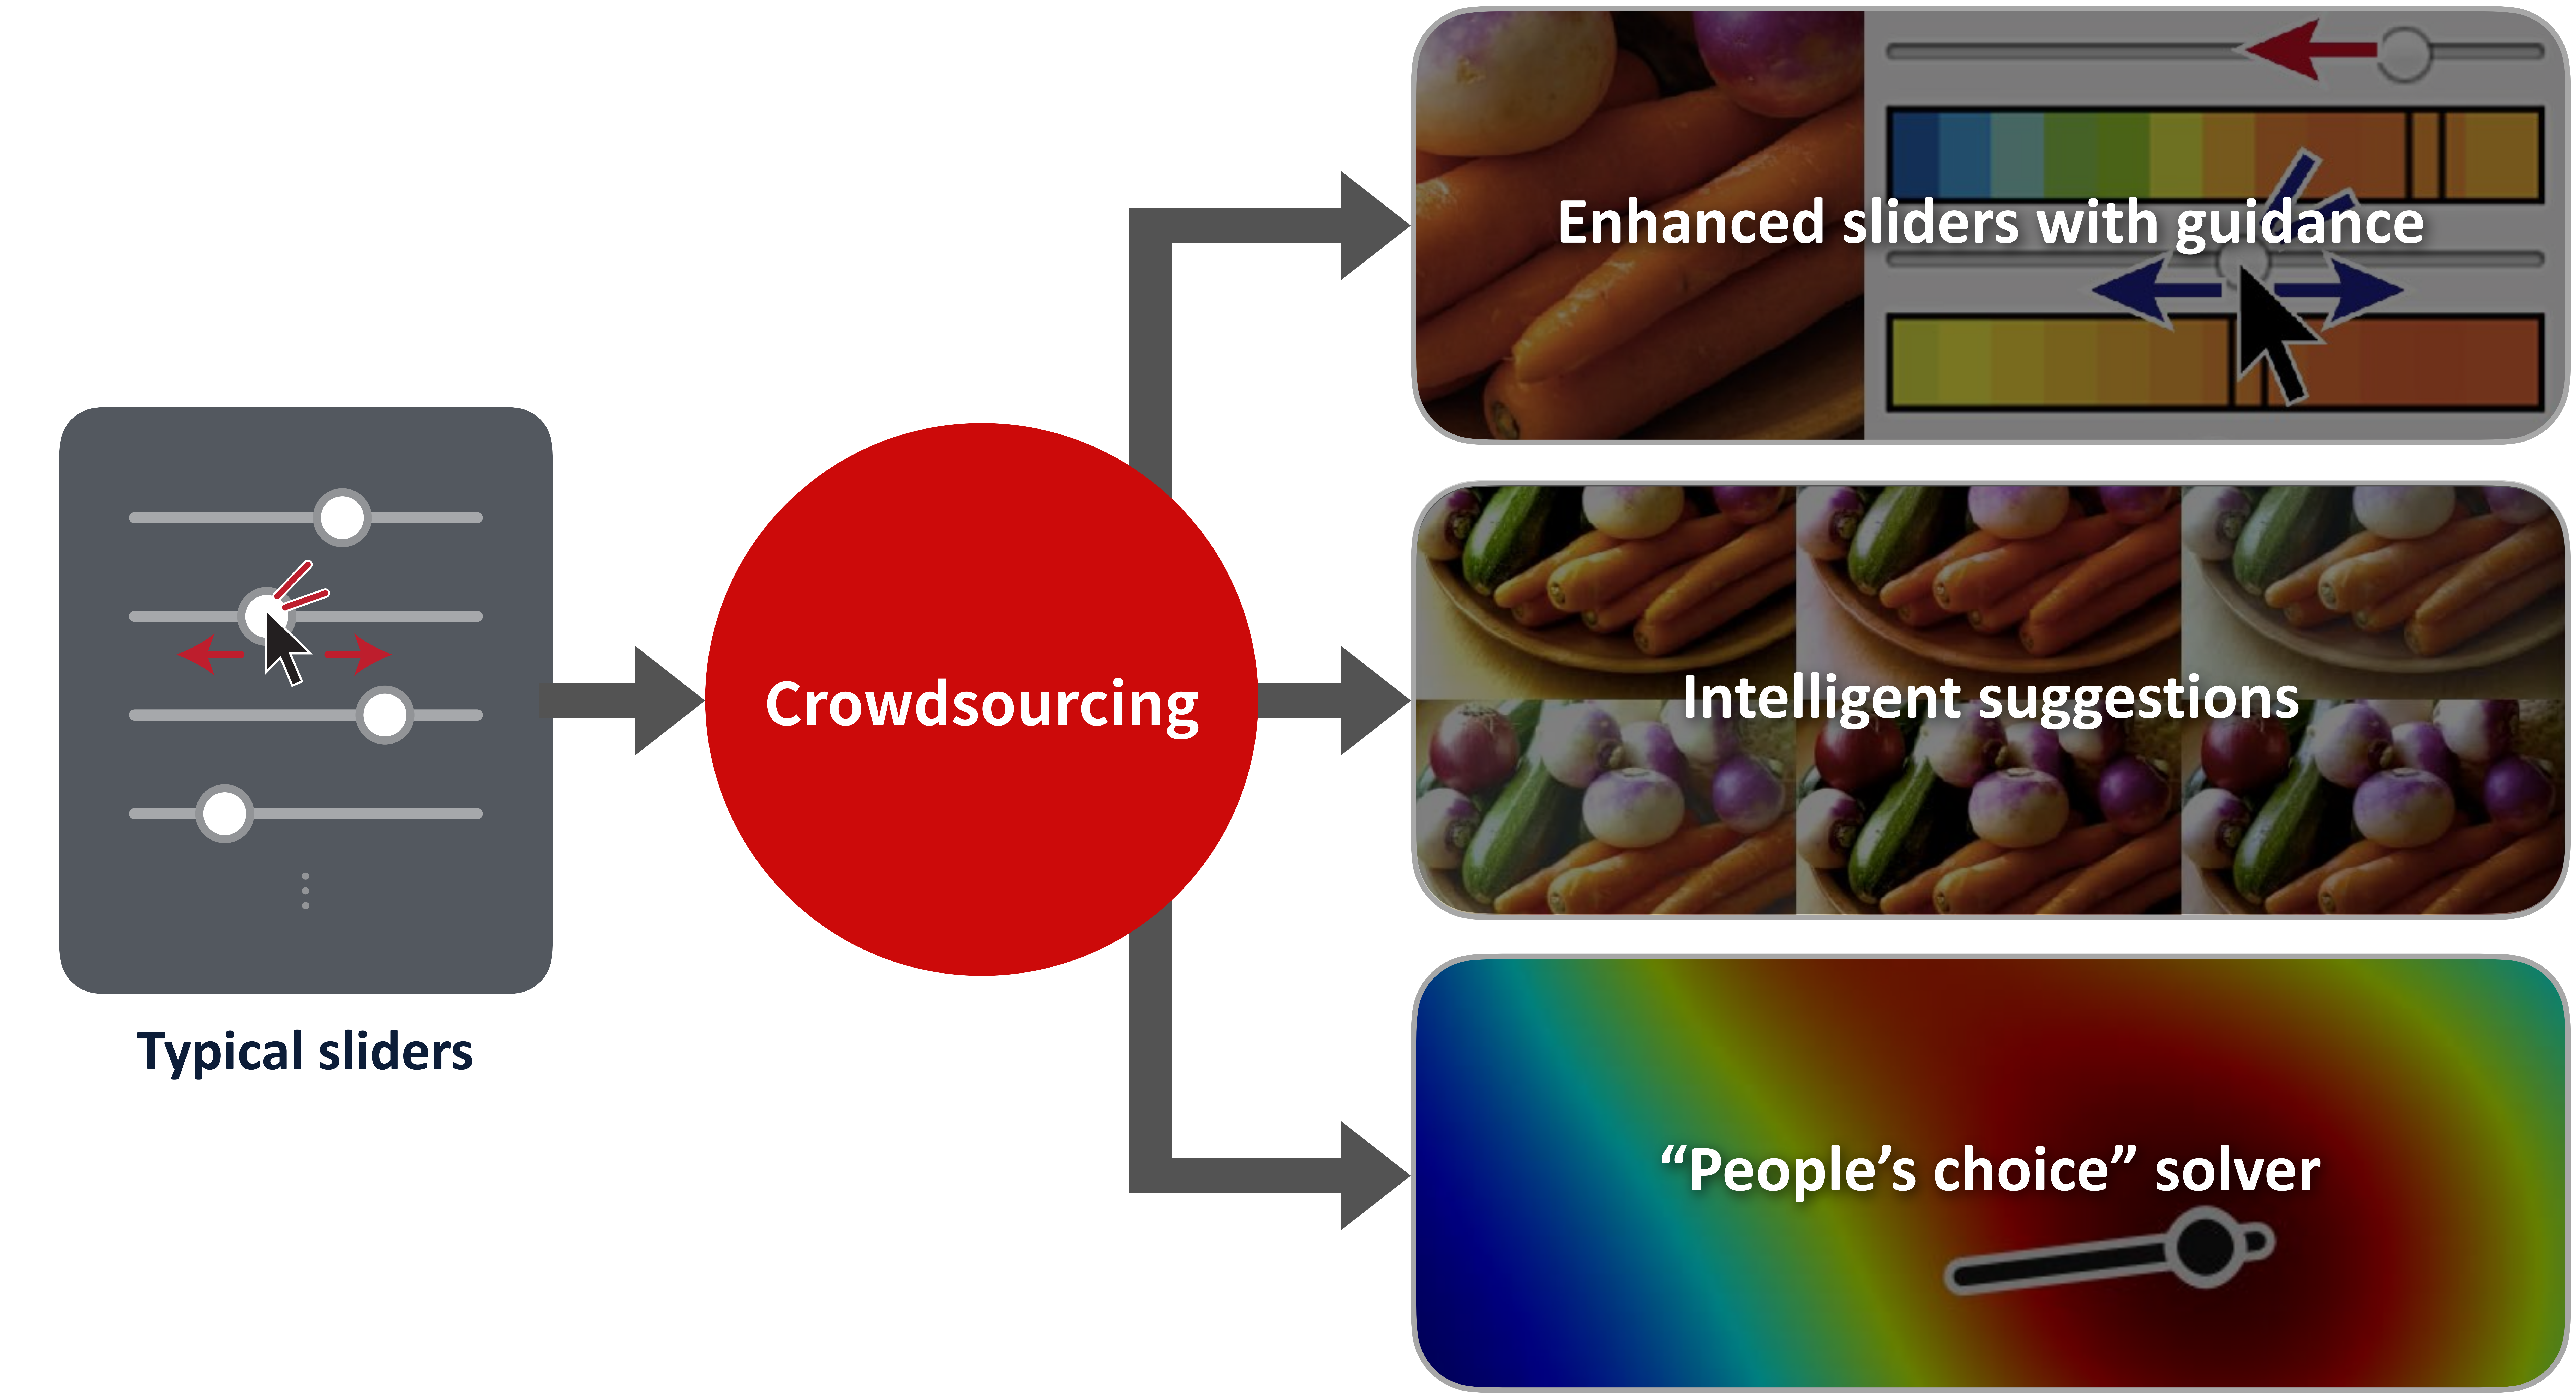
Typical sliders
Crowdsourcing
Enhanced sliders with guidance
Intelligent suggestions
"People's choice" solver

Basics

Crowdsourcing and Related Concepts

Concept: Crowdsourcing

- Jeff Howe (an author of WIRED) defined the term, crowdsourcing, in **2006**

“the act of taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call”

- This term itself is not related to computer science or software engineering

Jeff Howe. 2006. Crowdsourcing: A Definition. https://crowdsourcing.typepad.com/cs/2006/06/crowdsourcing_a.html

Concept: Crowdsourcing

Possible Styles

- **Microtask** (e.g., Amazon Mechanical Turk)
- Expert (e.g., Upwork)
- Volunteering (e.g., open calls in citizen science projects)
- … etc.

In the rest of this course, we assume the **microtask**-based crowdsourcing since it has good **reproducibility, stability, and scalability**, and we can use it **on demand**.

Concept: Human Computation

- Luis von Ahn (a pioneer of human computation) [2005] described it as

> "a paradigm for utilizing human processing power to solve problems that computers cannot yet solve"

- How to obtain **human processing power**
  - Gamification (FoldIt, …)
  - Embedding into existing tasks (reCAPTCHA, …)
  - **Microtask-based crowdsourcing** (next slide)

[von Ahn, 2005] Luis von Ahn. 2005. Human Computation. PhD thesis, Carnegie Mellon University, CMU-CS-05-193. http://reports-archive.adm.cs.cmu.edu/anon/2005/CMU-CS-05-193.pdf

Concept: Crowdsourced Human Computation

- Consider crowd workers as **intelligent processing power**
- Ask crowds to perform microtasks in the manner of function calls

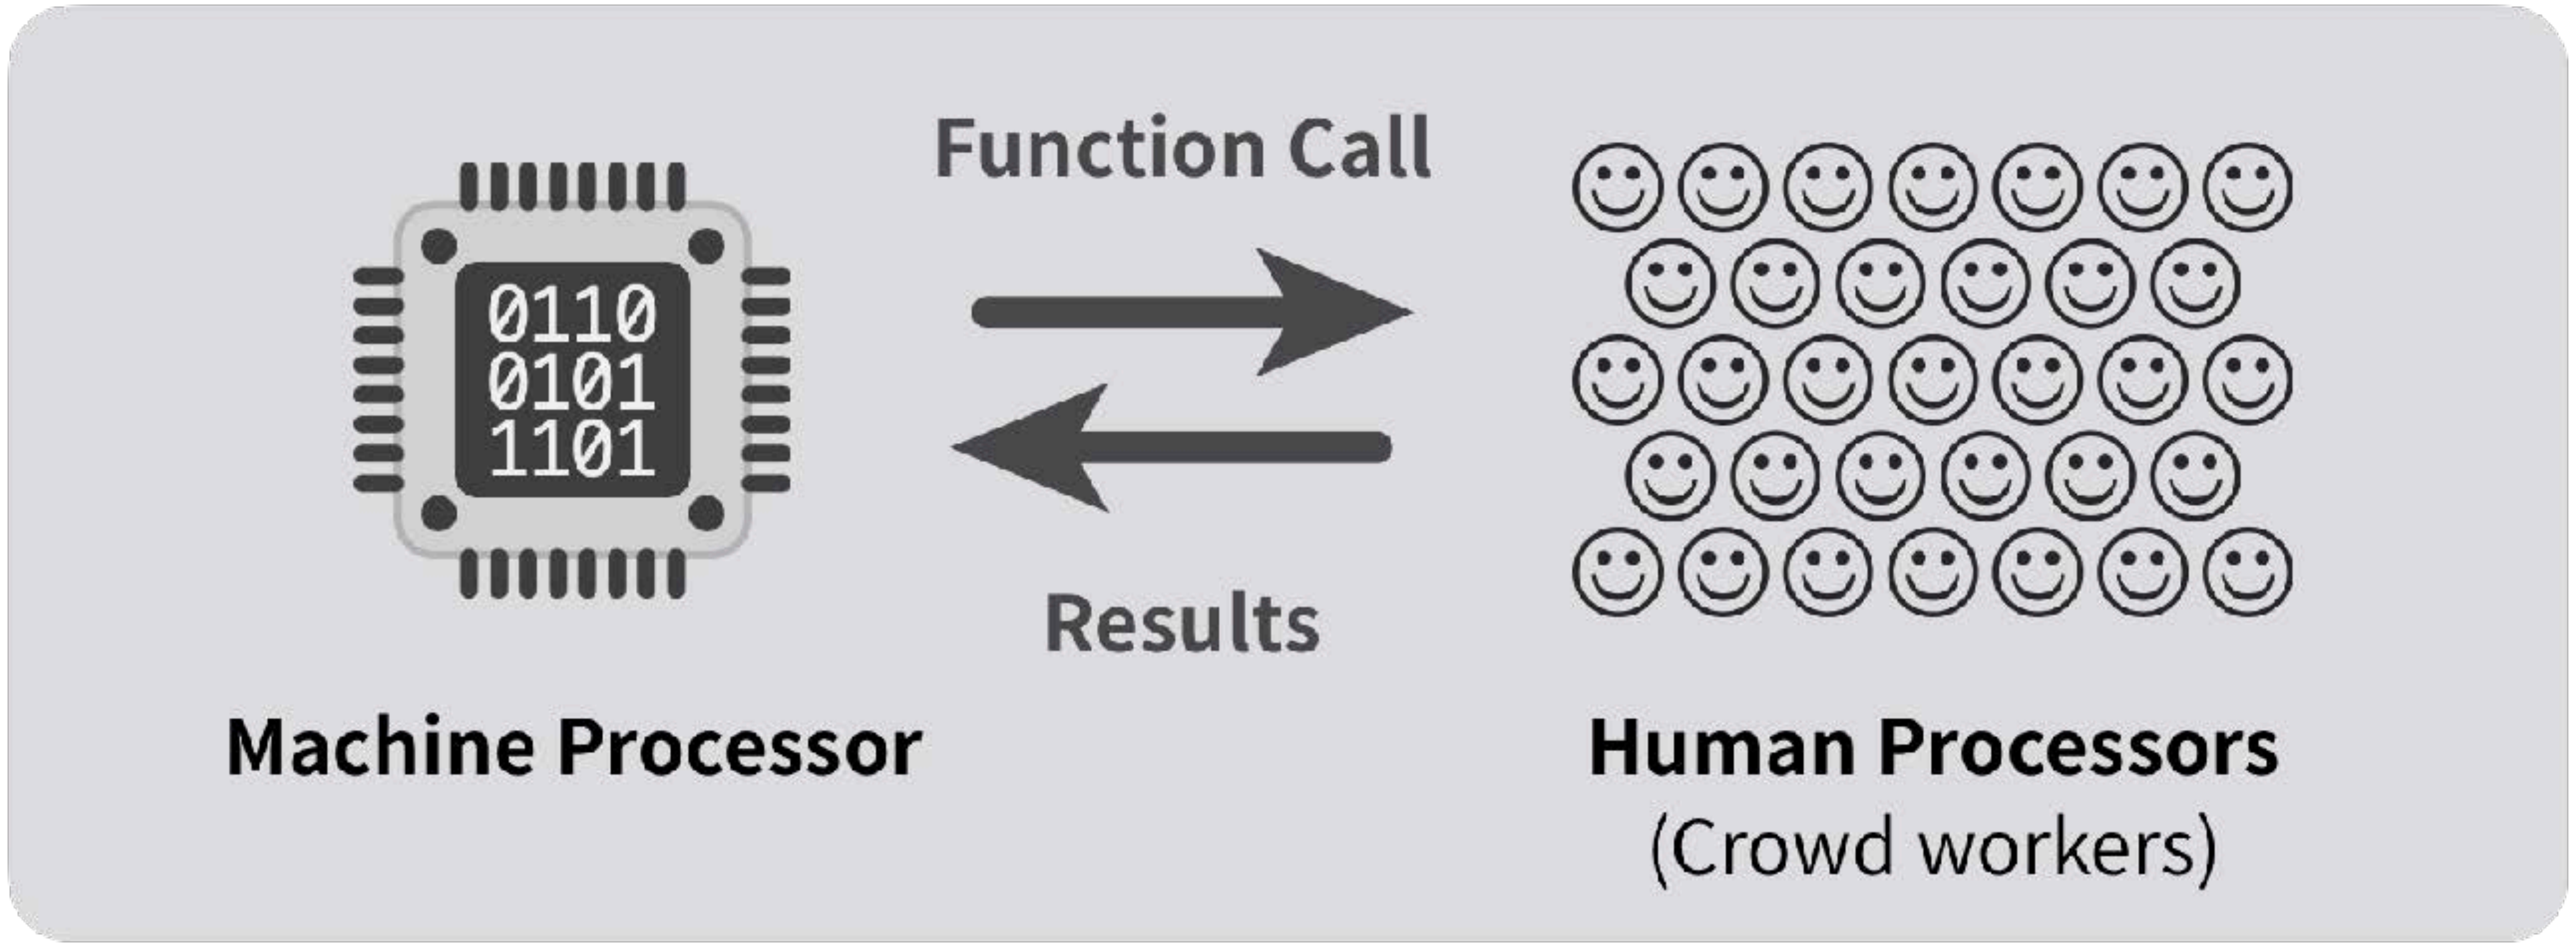

Concept: Crowdsourced Human Computation

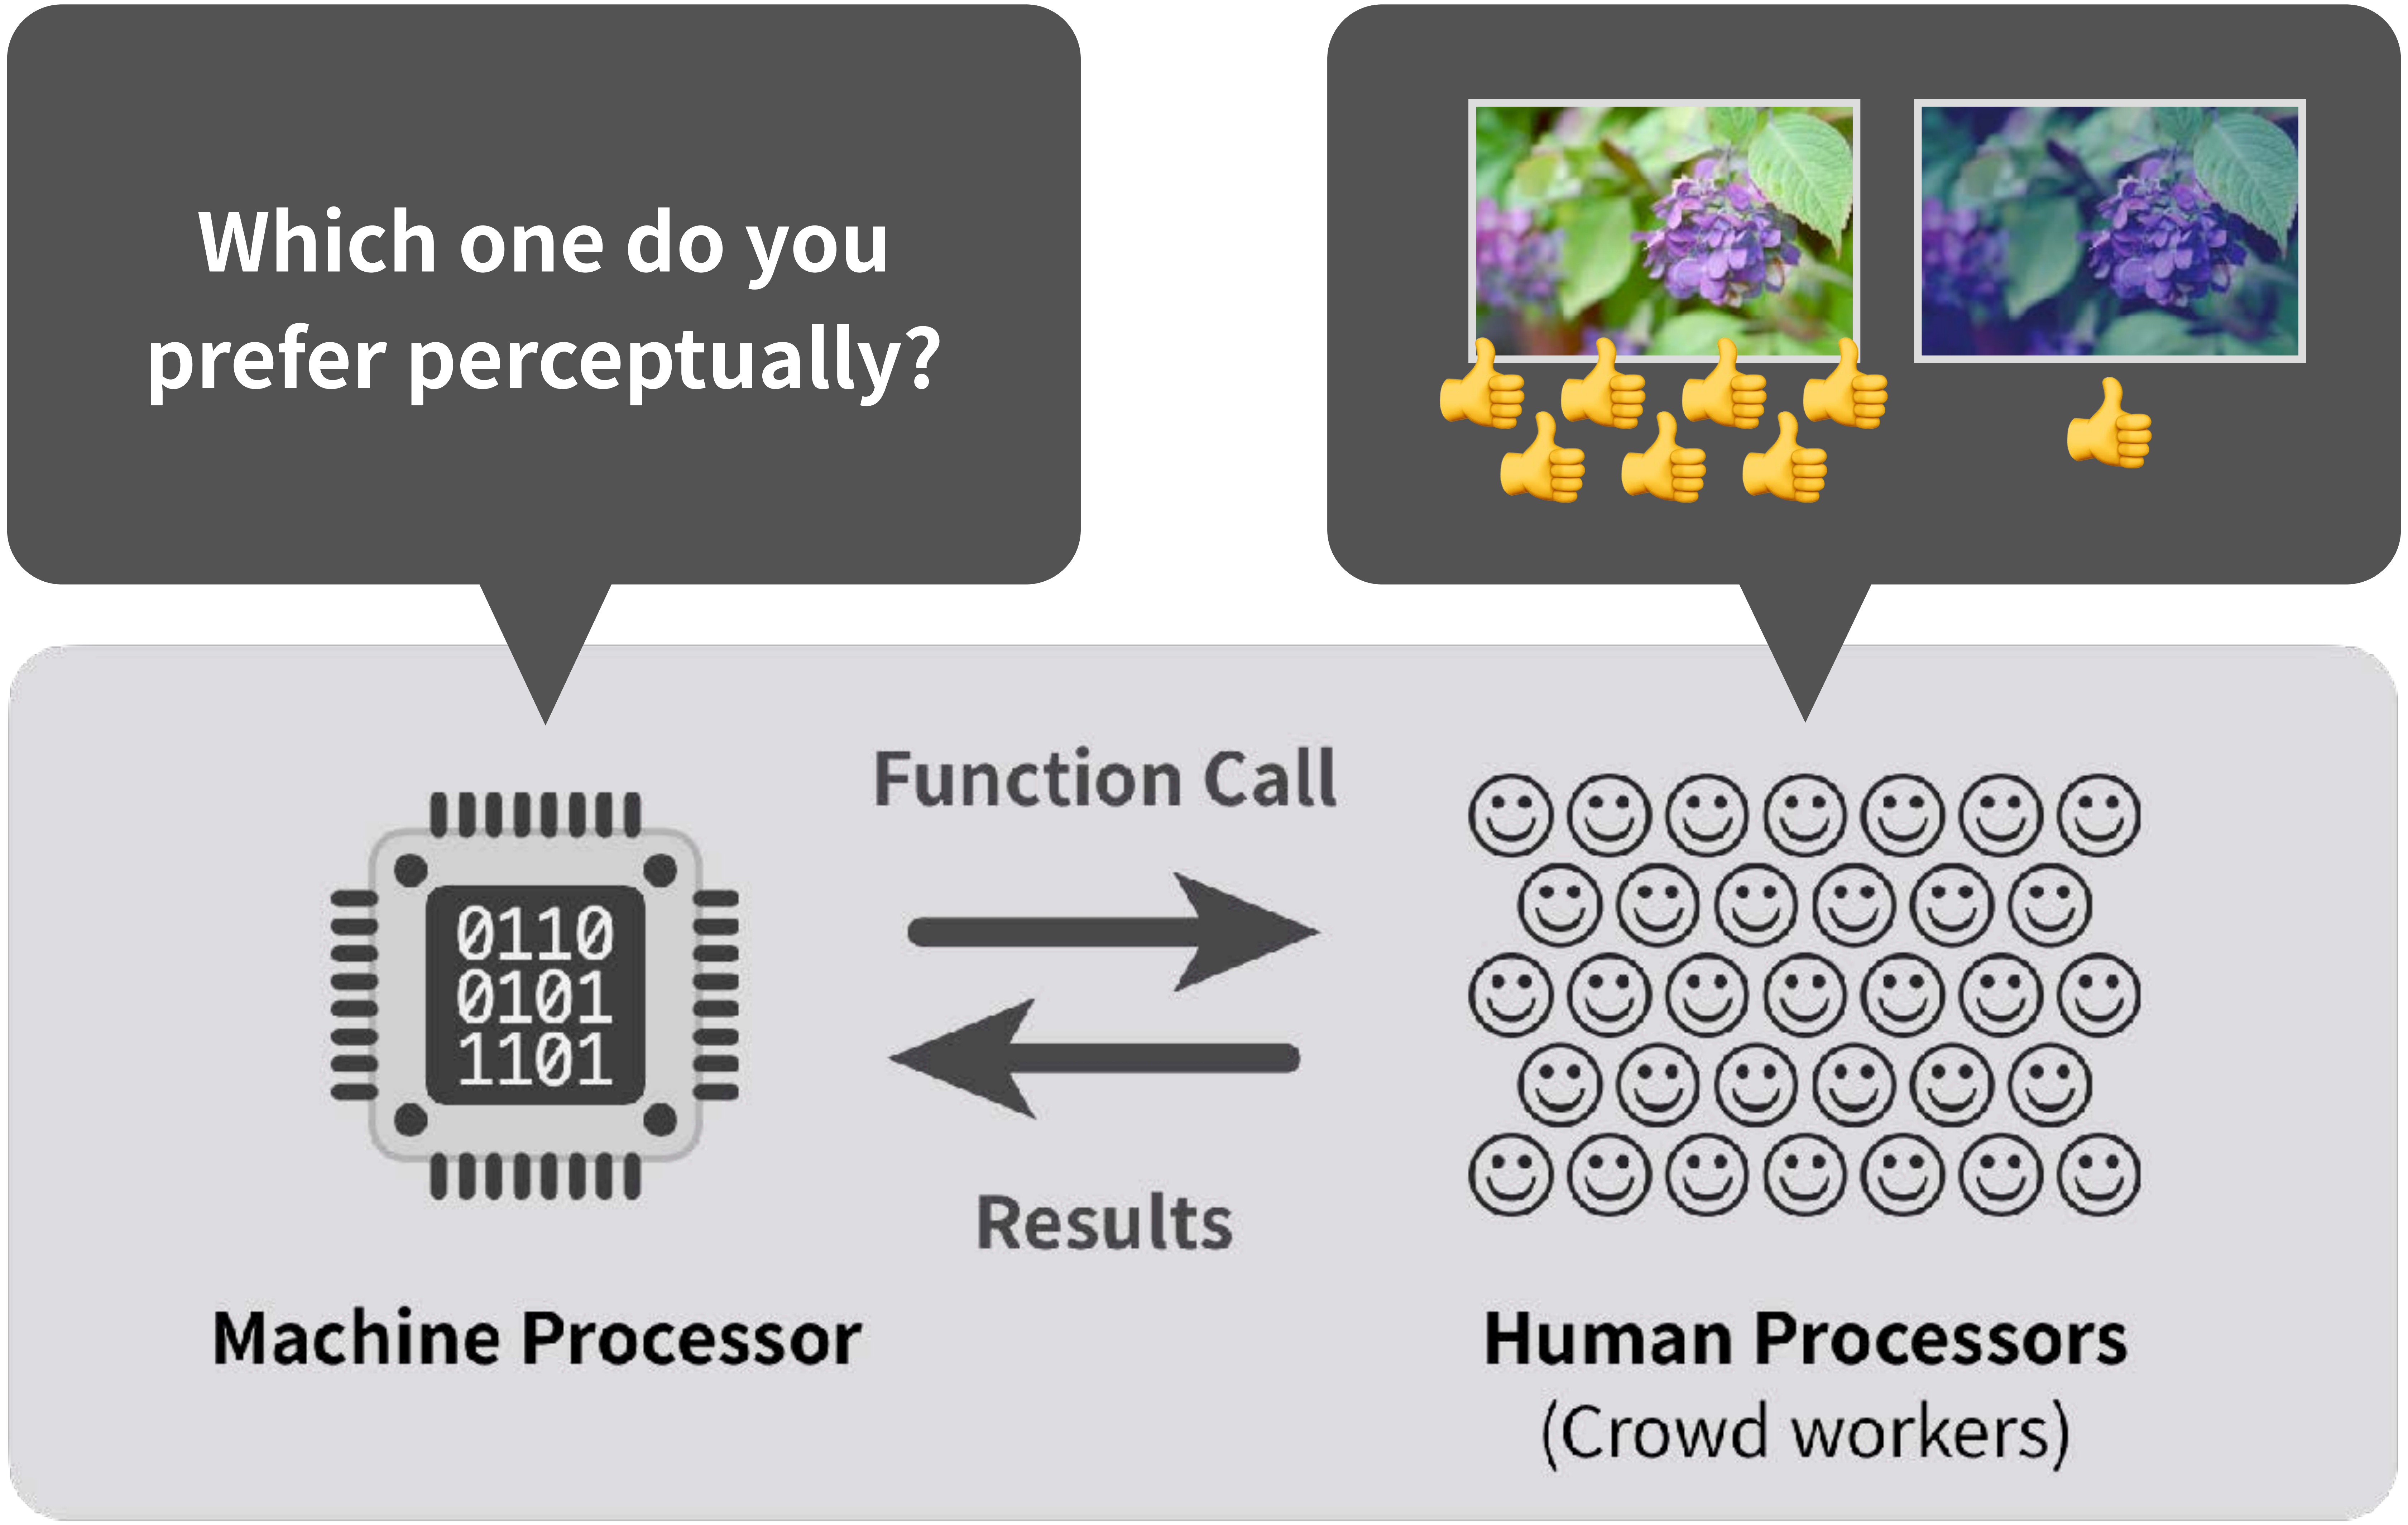

Crowdsourced Human Computation: Example

DEPTH-LAYERS(image I, sentinel queries S)

Segment I into regions (using mean-shift and SLIC)
Insert all pairs of neighboring regions into Q
loop in parallel until each pair has been visited N times
    Gather K random pairs from Q
    Gather M random pairs from S
    for each pair: Build the visual query & Duplicate it
    Mix the $2K + 2M$ queries
    $results$ = send all queries to an HP
    if $average(consistency(results)) \geq 0.75$ and
        $average(sentinel(results)) \geq 0.75$
        for each pair
            Add consistent results to the list of votes
            Increment #visited
for each pair of neighboring regions
    $final\_result = majority$(list of votes)
Solve the Laplace equation to construct a depth map

Crowds-in-the-loop algorithm

Input image

Depth map by crowds

[Gingold+, TOG (2012)] Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. 2012. Micro perceptual human computation for visual tasks. ACM Trans. Graph. 31, 5, 119:1–119:12 (2012). https://doi.org/10.1145/2231816.2231817

Crowdsourced Human Computation: Example

```
DEPTH-LAYERS(image I, sentinel queries S)
Segment I into regions (using mean-shift and SLIC)
Insert all pairs of neighboring regions into Q
loop in parallel until each pair has been visited N times
    Gather K random pairs from Q
    Gather M random pairs from S
    for each pair: Build the visual query & Duplicate it
    Mix the 2K + 2M queries
    results = send all queries to an HP
    if average(consistency(results)) ≥ 0.75 and
      average(sentinel(results)) ≥ 0.75
        for each pair
            Add consistent results to the list of votes
            Increment #visited
for each pair of neighboring regions
    final_result = majority(list of votes)
Solve the Laplace equation to construct a depth map
```

$results$ = send all queries to an HP

HP: human processor

Crowds-in-the-loop algorithm

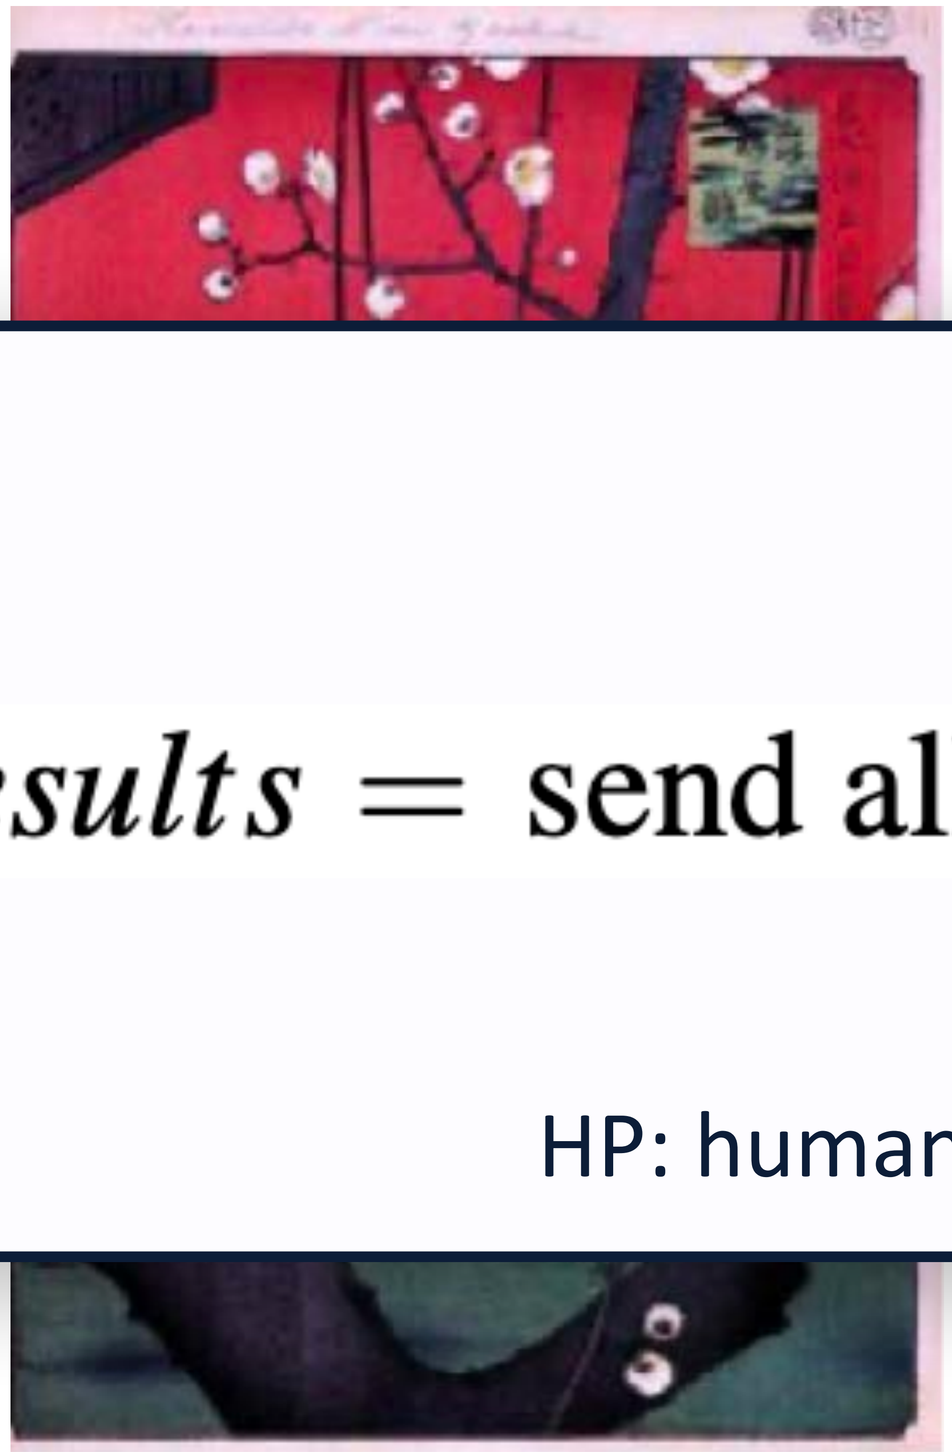

Input image

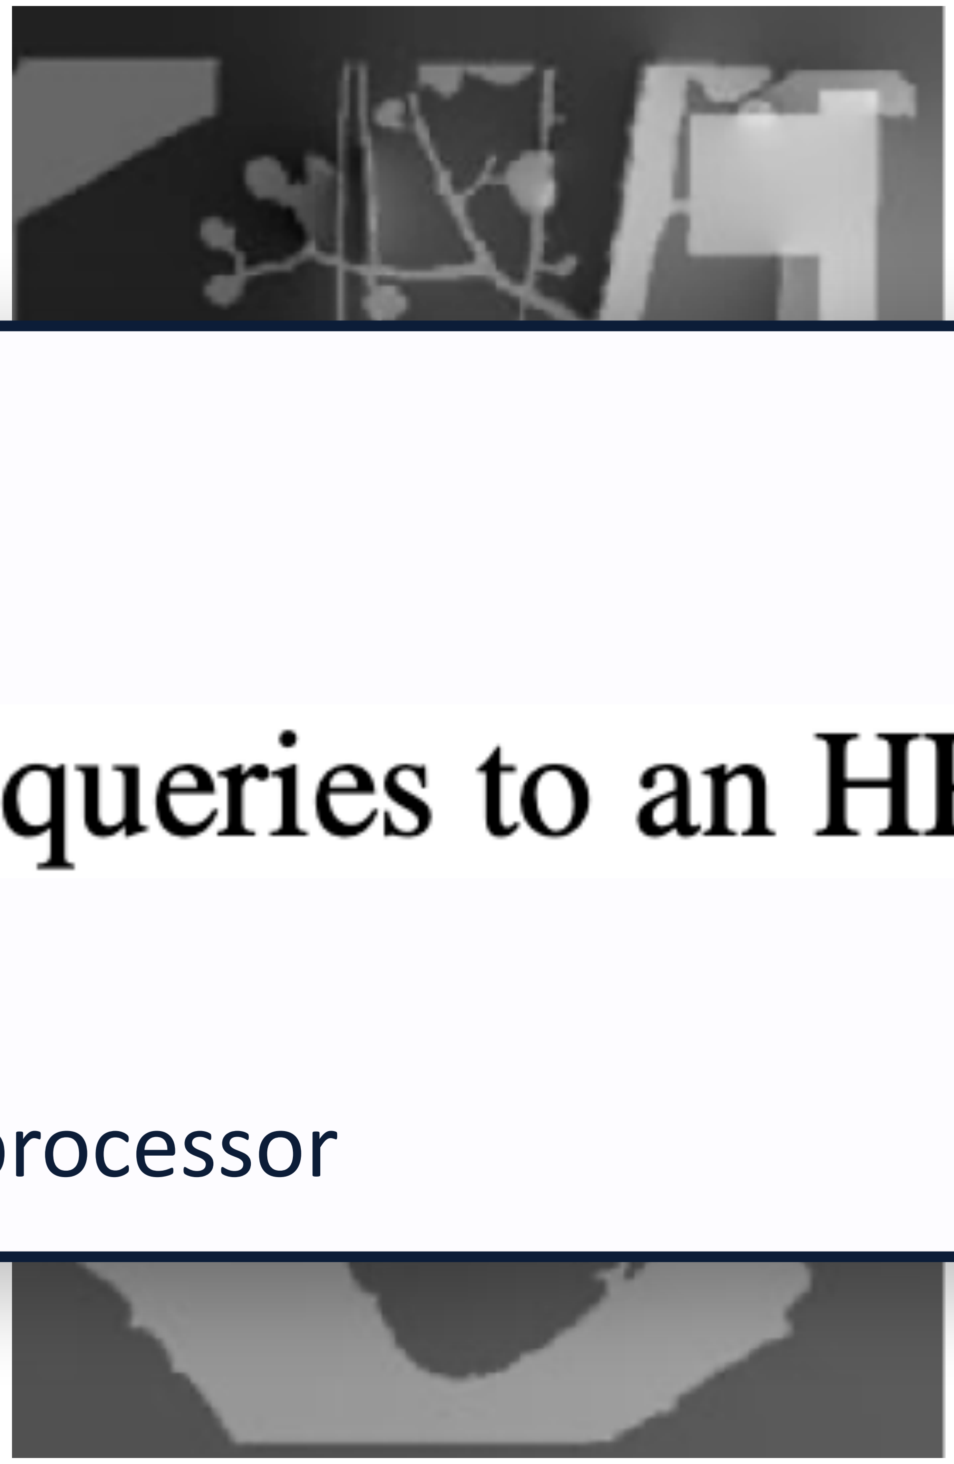

Depth map by crowds

[Gingold+, TOG (2012)] Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. 2012. Micro perceptual human computation for visual tasks. ACM Trans. Graph. 31, 5, 119:1–119:12 (2012). https://doi.org/10.1145/2231816.2231817

Formulation

Perceptual Feedback from Crowds

Problem Definition from Mathematical Viewpoint

- Suppose that we have n sliders to adjust
- Let $\mathscr{X} = [0,1]^n$ be the search space and $\mathbf{x} \in \mathscr{X}$ be a set of n parameter values
- Let $g : \mathscr{X} \to \mathbb{R}$ be a perceptual preference function (= **goodness function**) which returns a goodness value
- We want to solve an optimization problem:

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

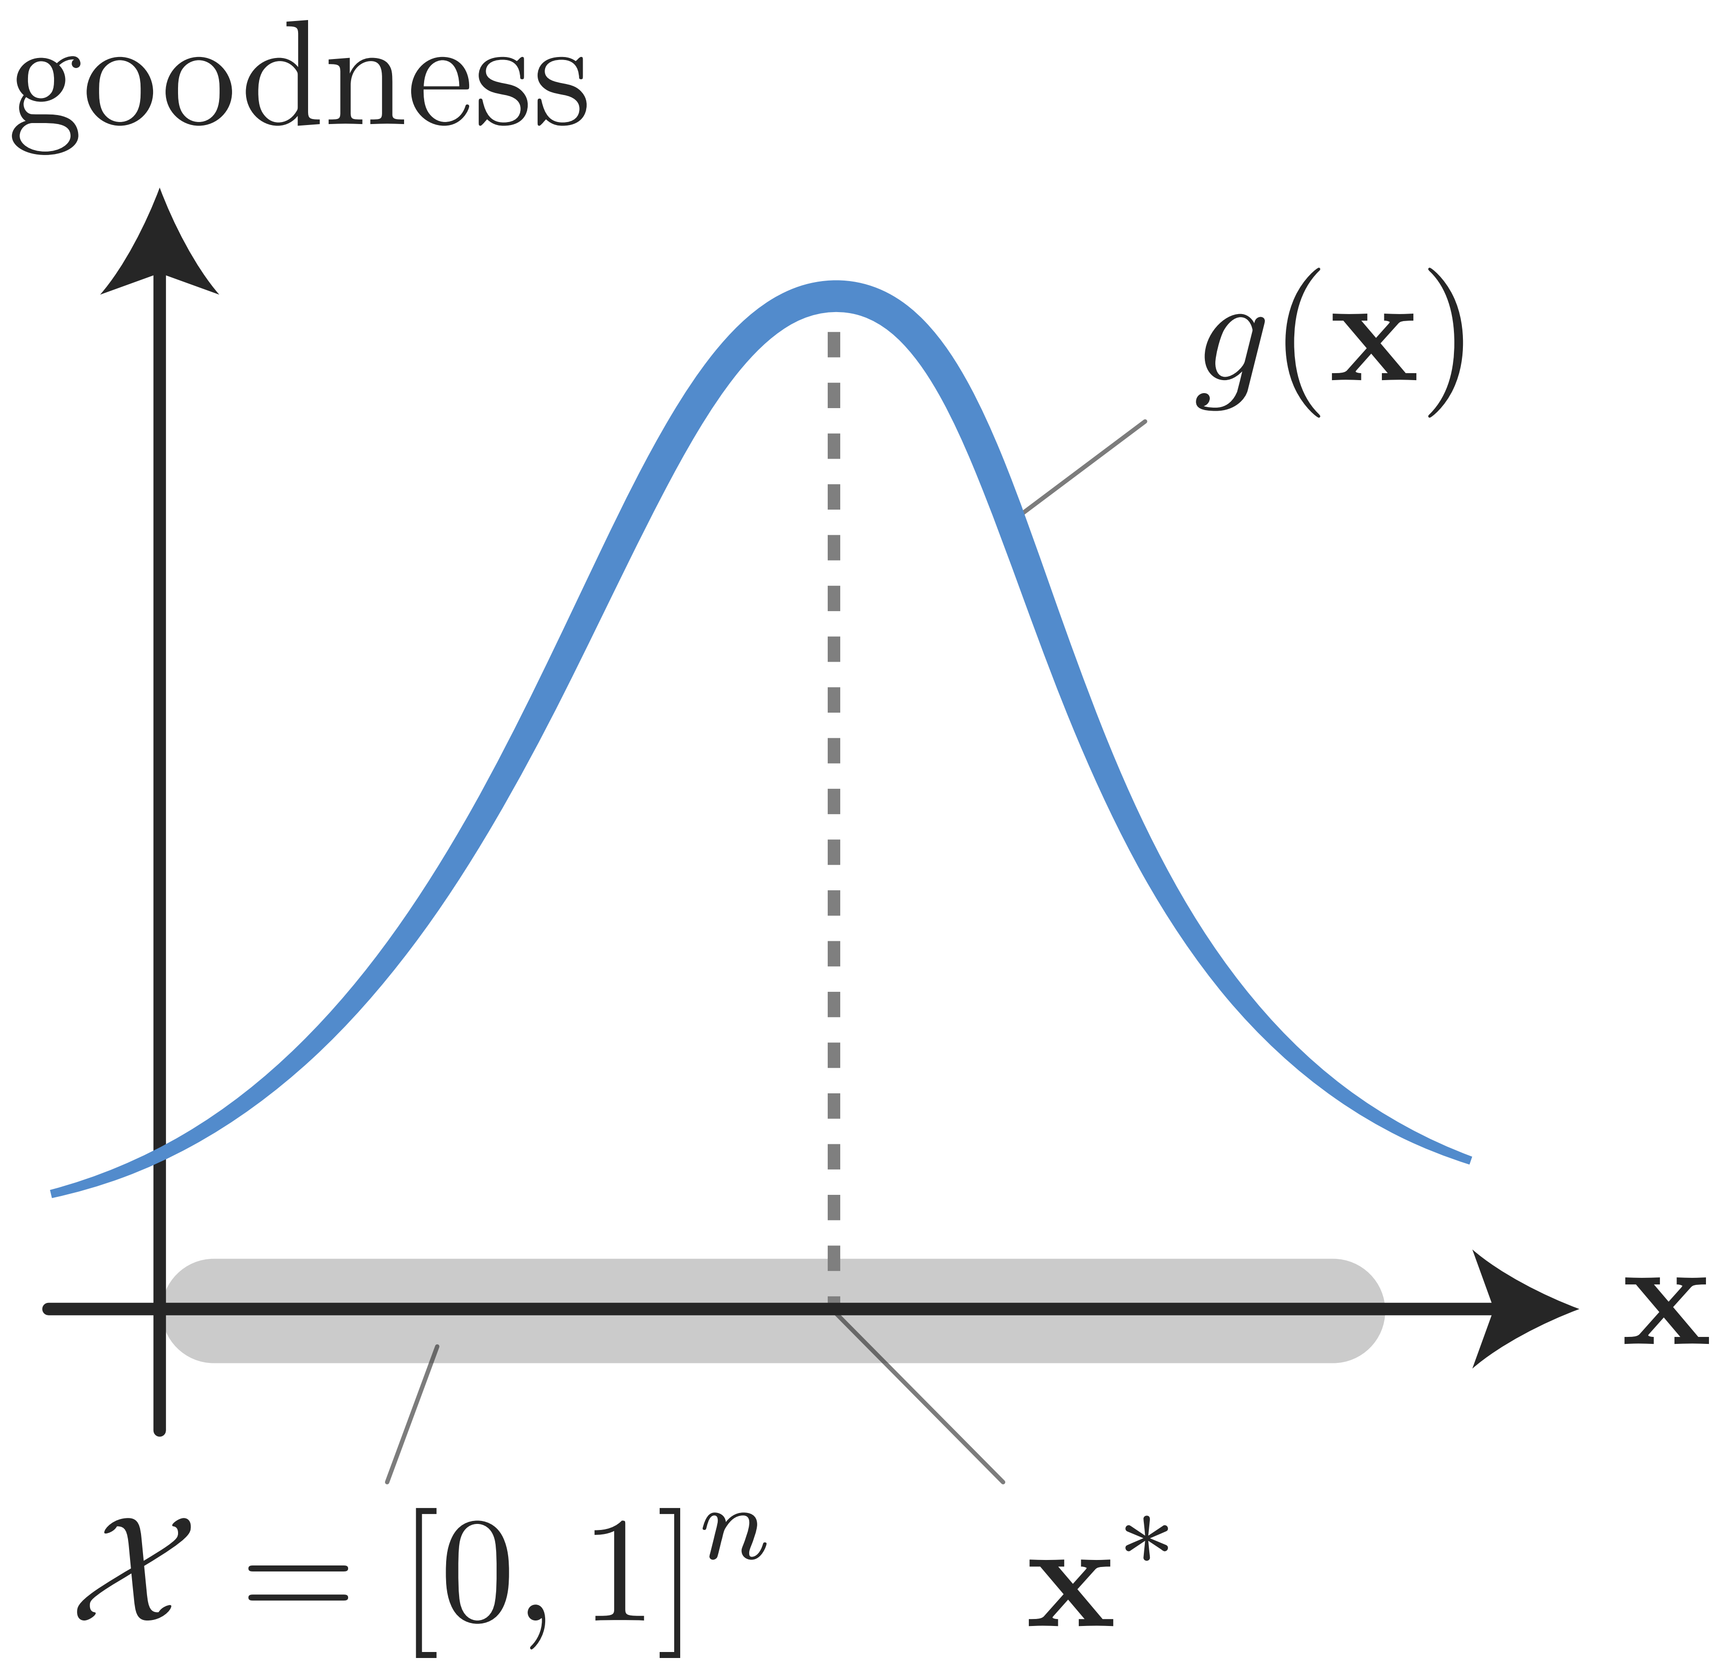

[Koyama+, Computational Interaction (2018)] Yuki Koyama and Takeo Igarashi. 2018. Computational Design with Crowds. In Computational Interaction (Eds. Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes), Oxford University Press, pp.153—184.

Interacting with Perceptual Functions

- **Absolute assessment should not be used:** Crowds cannot directly answer the function value reliably [Koyama+18]

[Koyama+, Computational Interaction (2018)] Yuki Koyama and Takeo Igarashi. 2018. Computational Design with Crowds. In Computational Interaction (Eds. A. Oulasvirta, P. O. Kristensson, X. Bi, and A. Howes), Oxford University Press, pp.153–184. https://arxiv.org/abs/2002.08657

Interacting with Perceptual Functions

- **Absolute assessment should not be used:** Crowds cannot directly answer the function value reliably [Koyama+18]
- **Relative assessment should be used:** Crowds can answer which option is better among two (or more) options

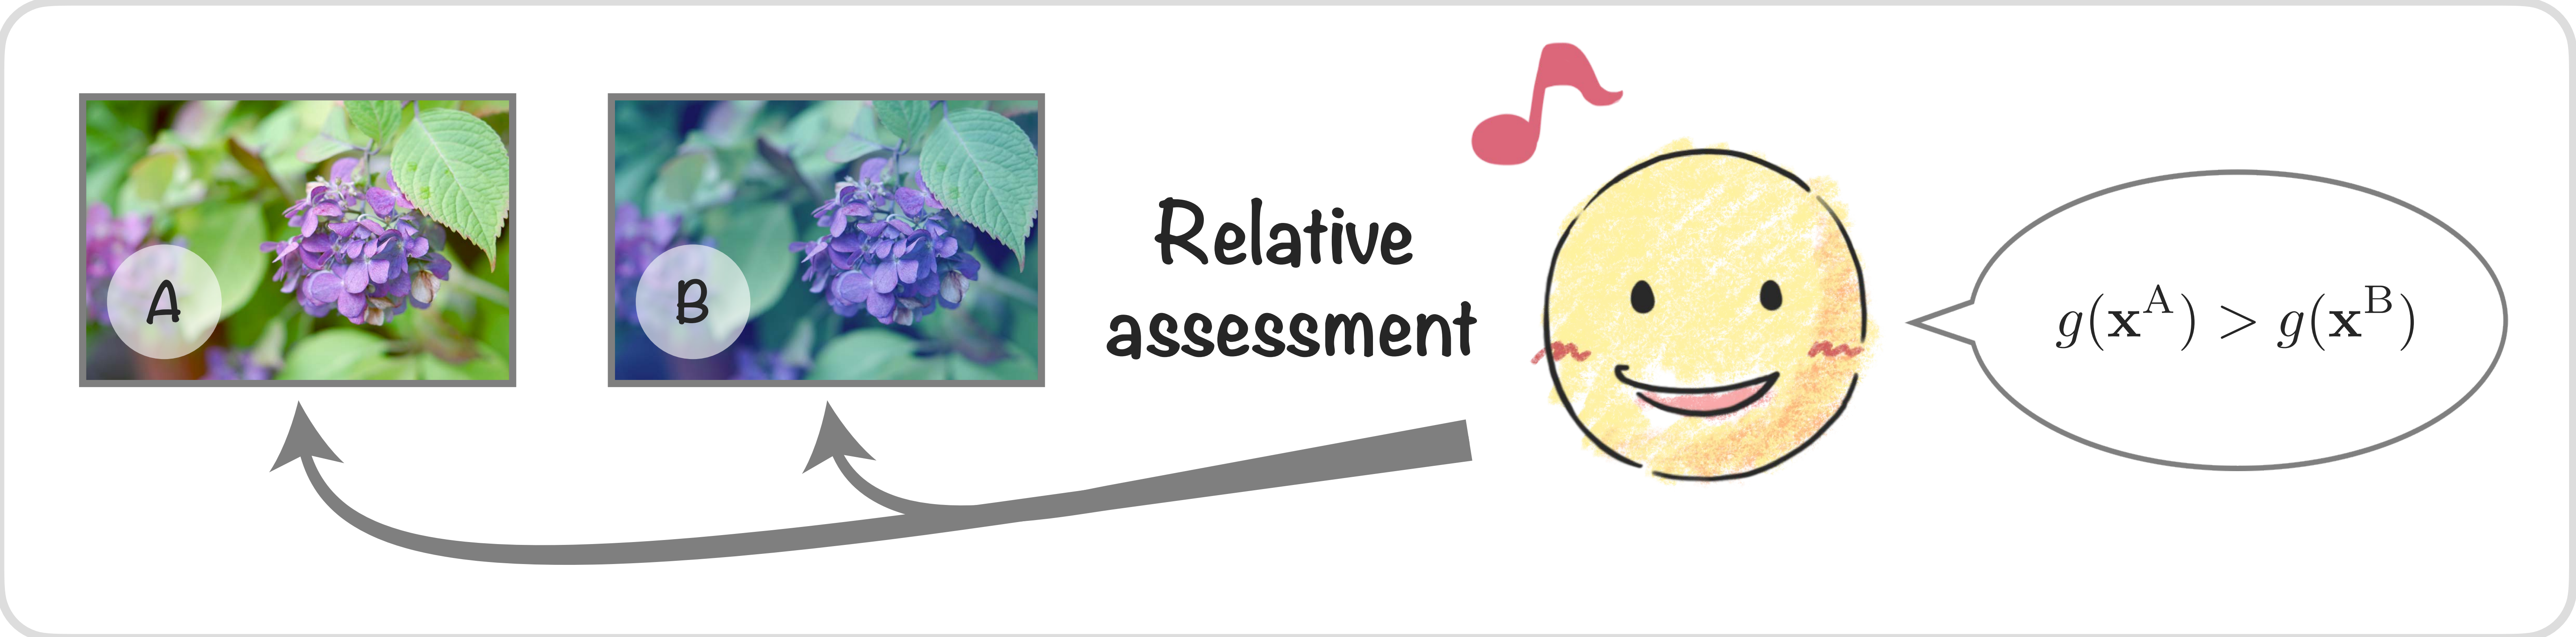

[Koyama+, Computational Interaction (2018)] Yuki Koyama and Takeo Igarashi. 2018. Computational Design with Crowds. In Computational Interaction (Eds. A. Oulasvirta, P. O. Kristensson, X. Bi, and A. Howes), Oxford University Press, pp.153–184. https://arxiv.org/abs/2002.08657

Intelligent Tools Case 1

Intelligent Suggestions and Sliders

[Koyama+, UIST 2014] Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. 2014. Crowd-powered parameter analysis for visual design exploration. In Proc. UIST '14. pp.65–74. https://doi.org/10.1145/2642918.2647386

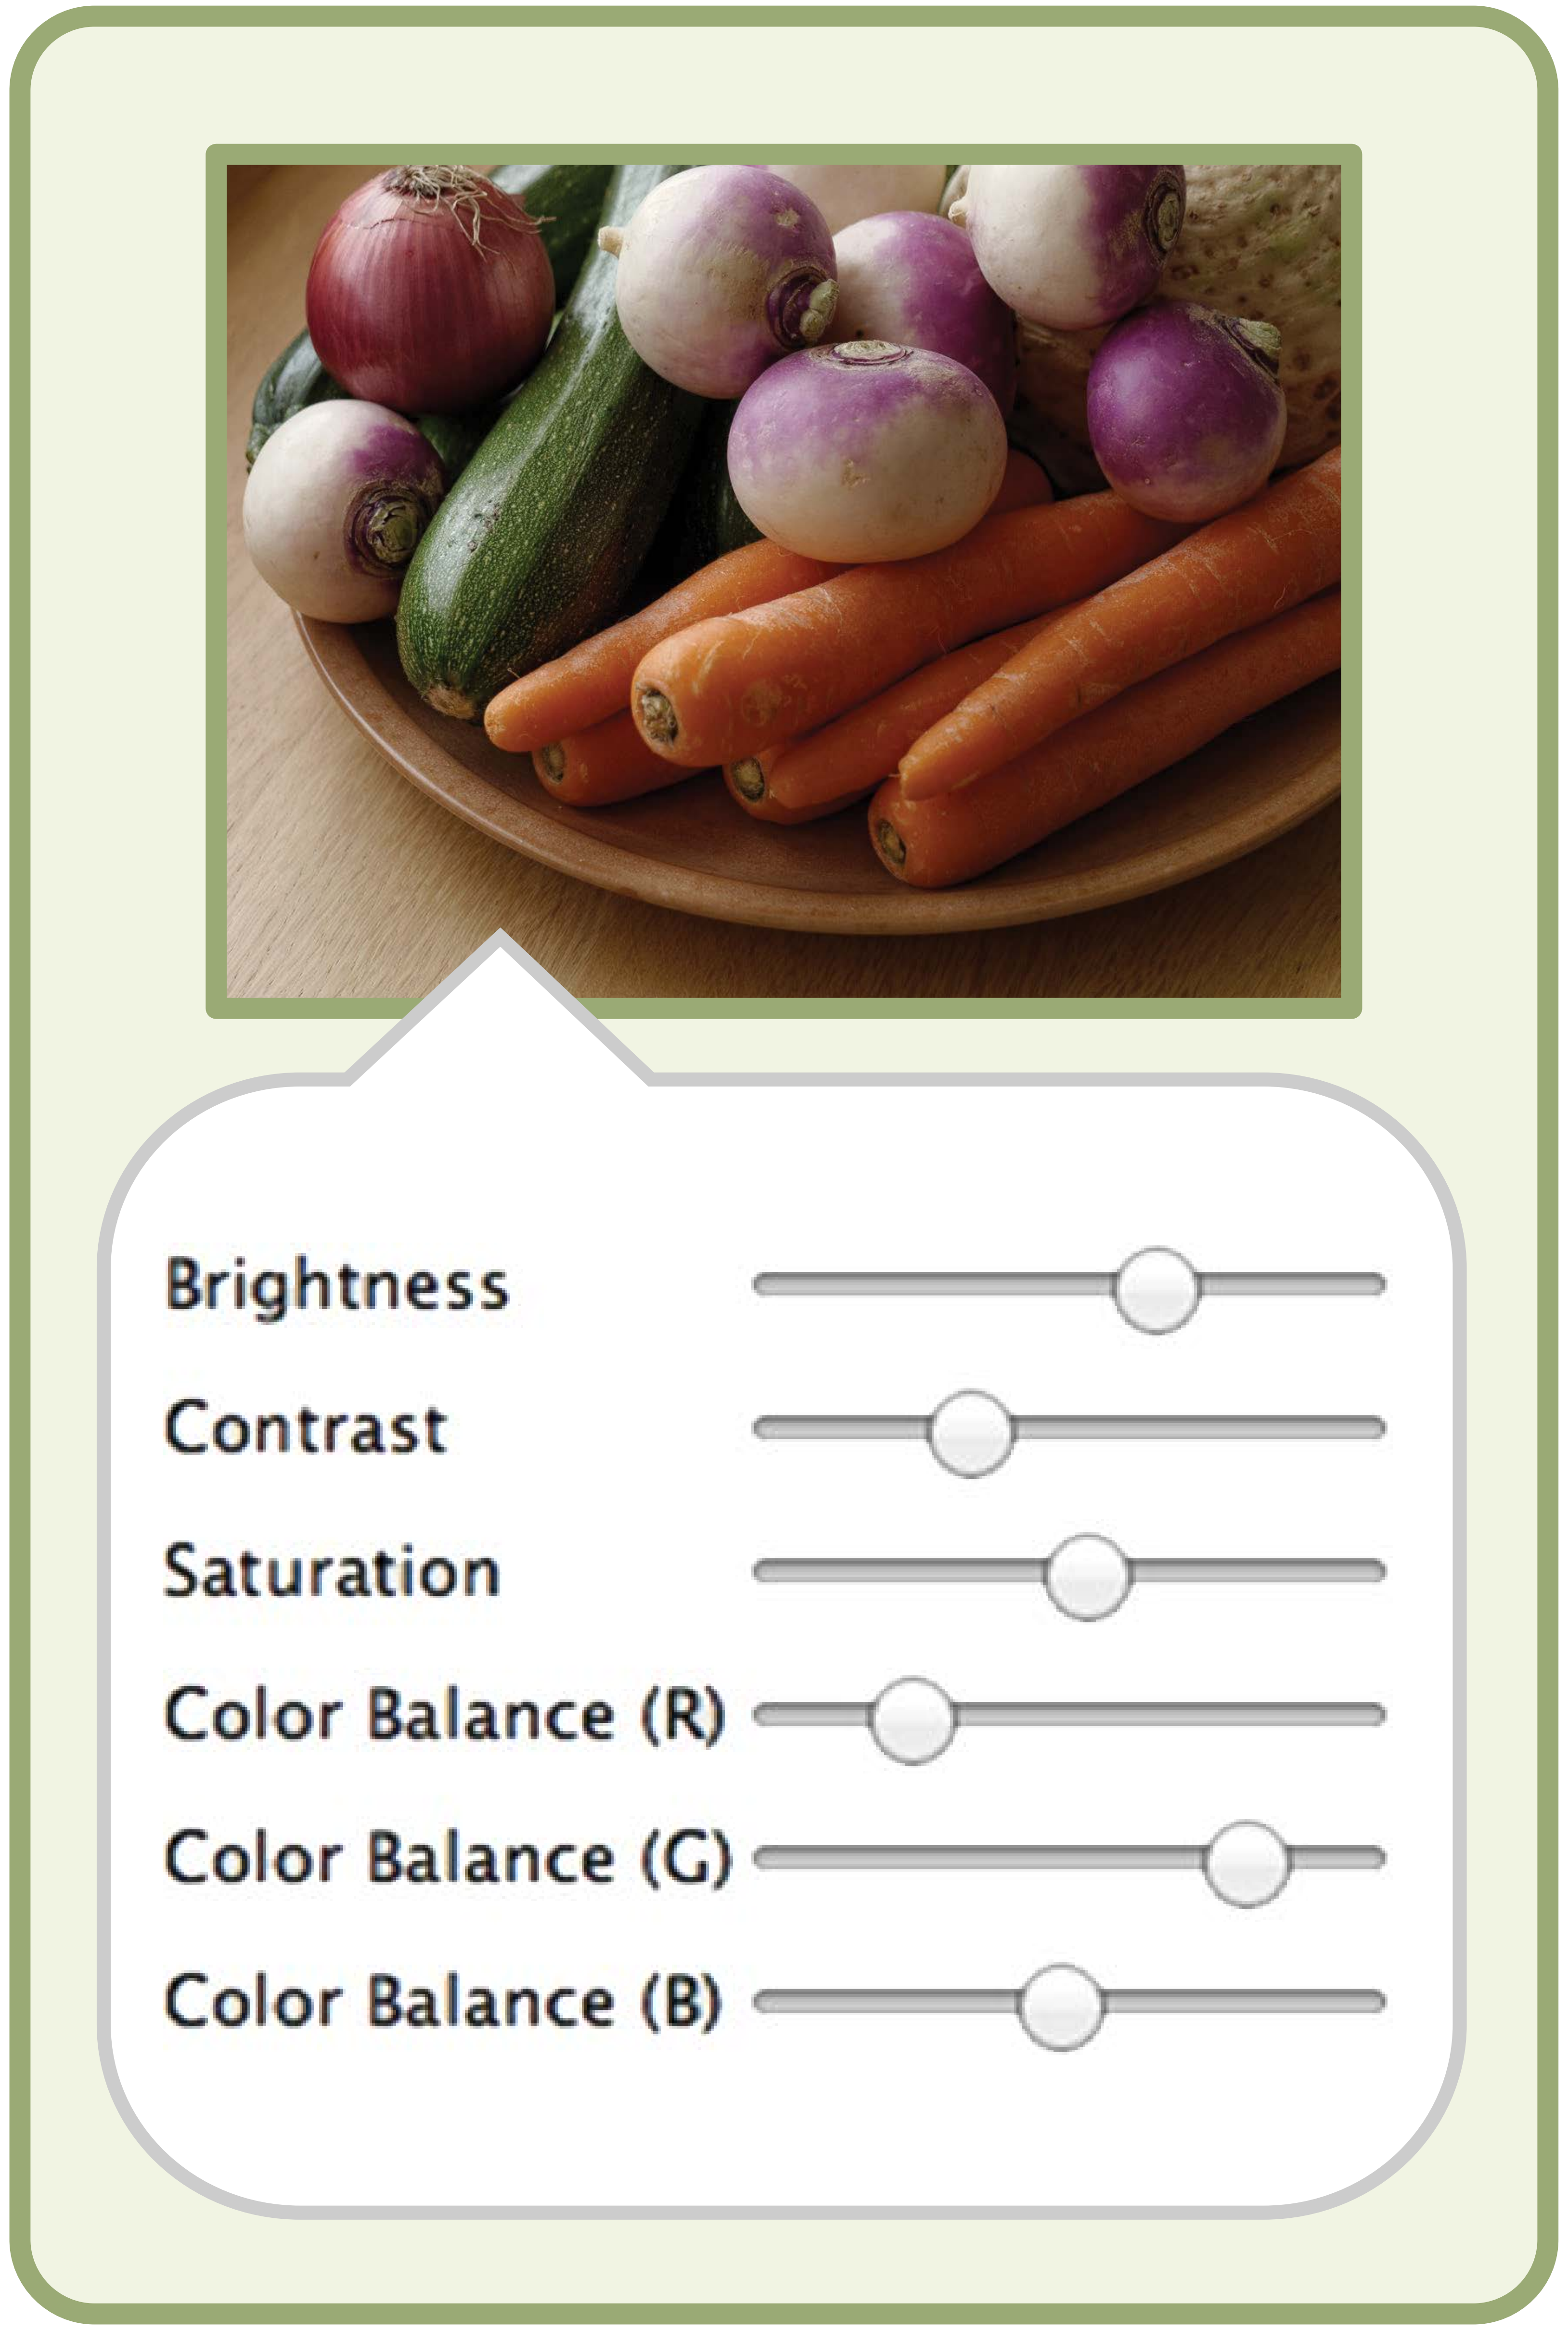

Target Parameters

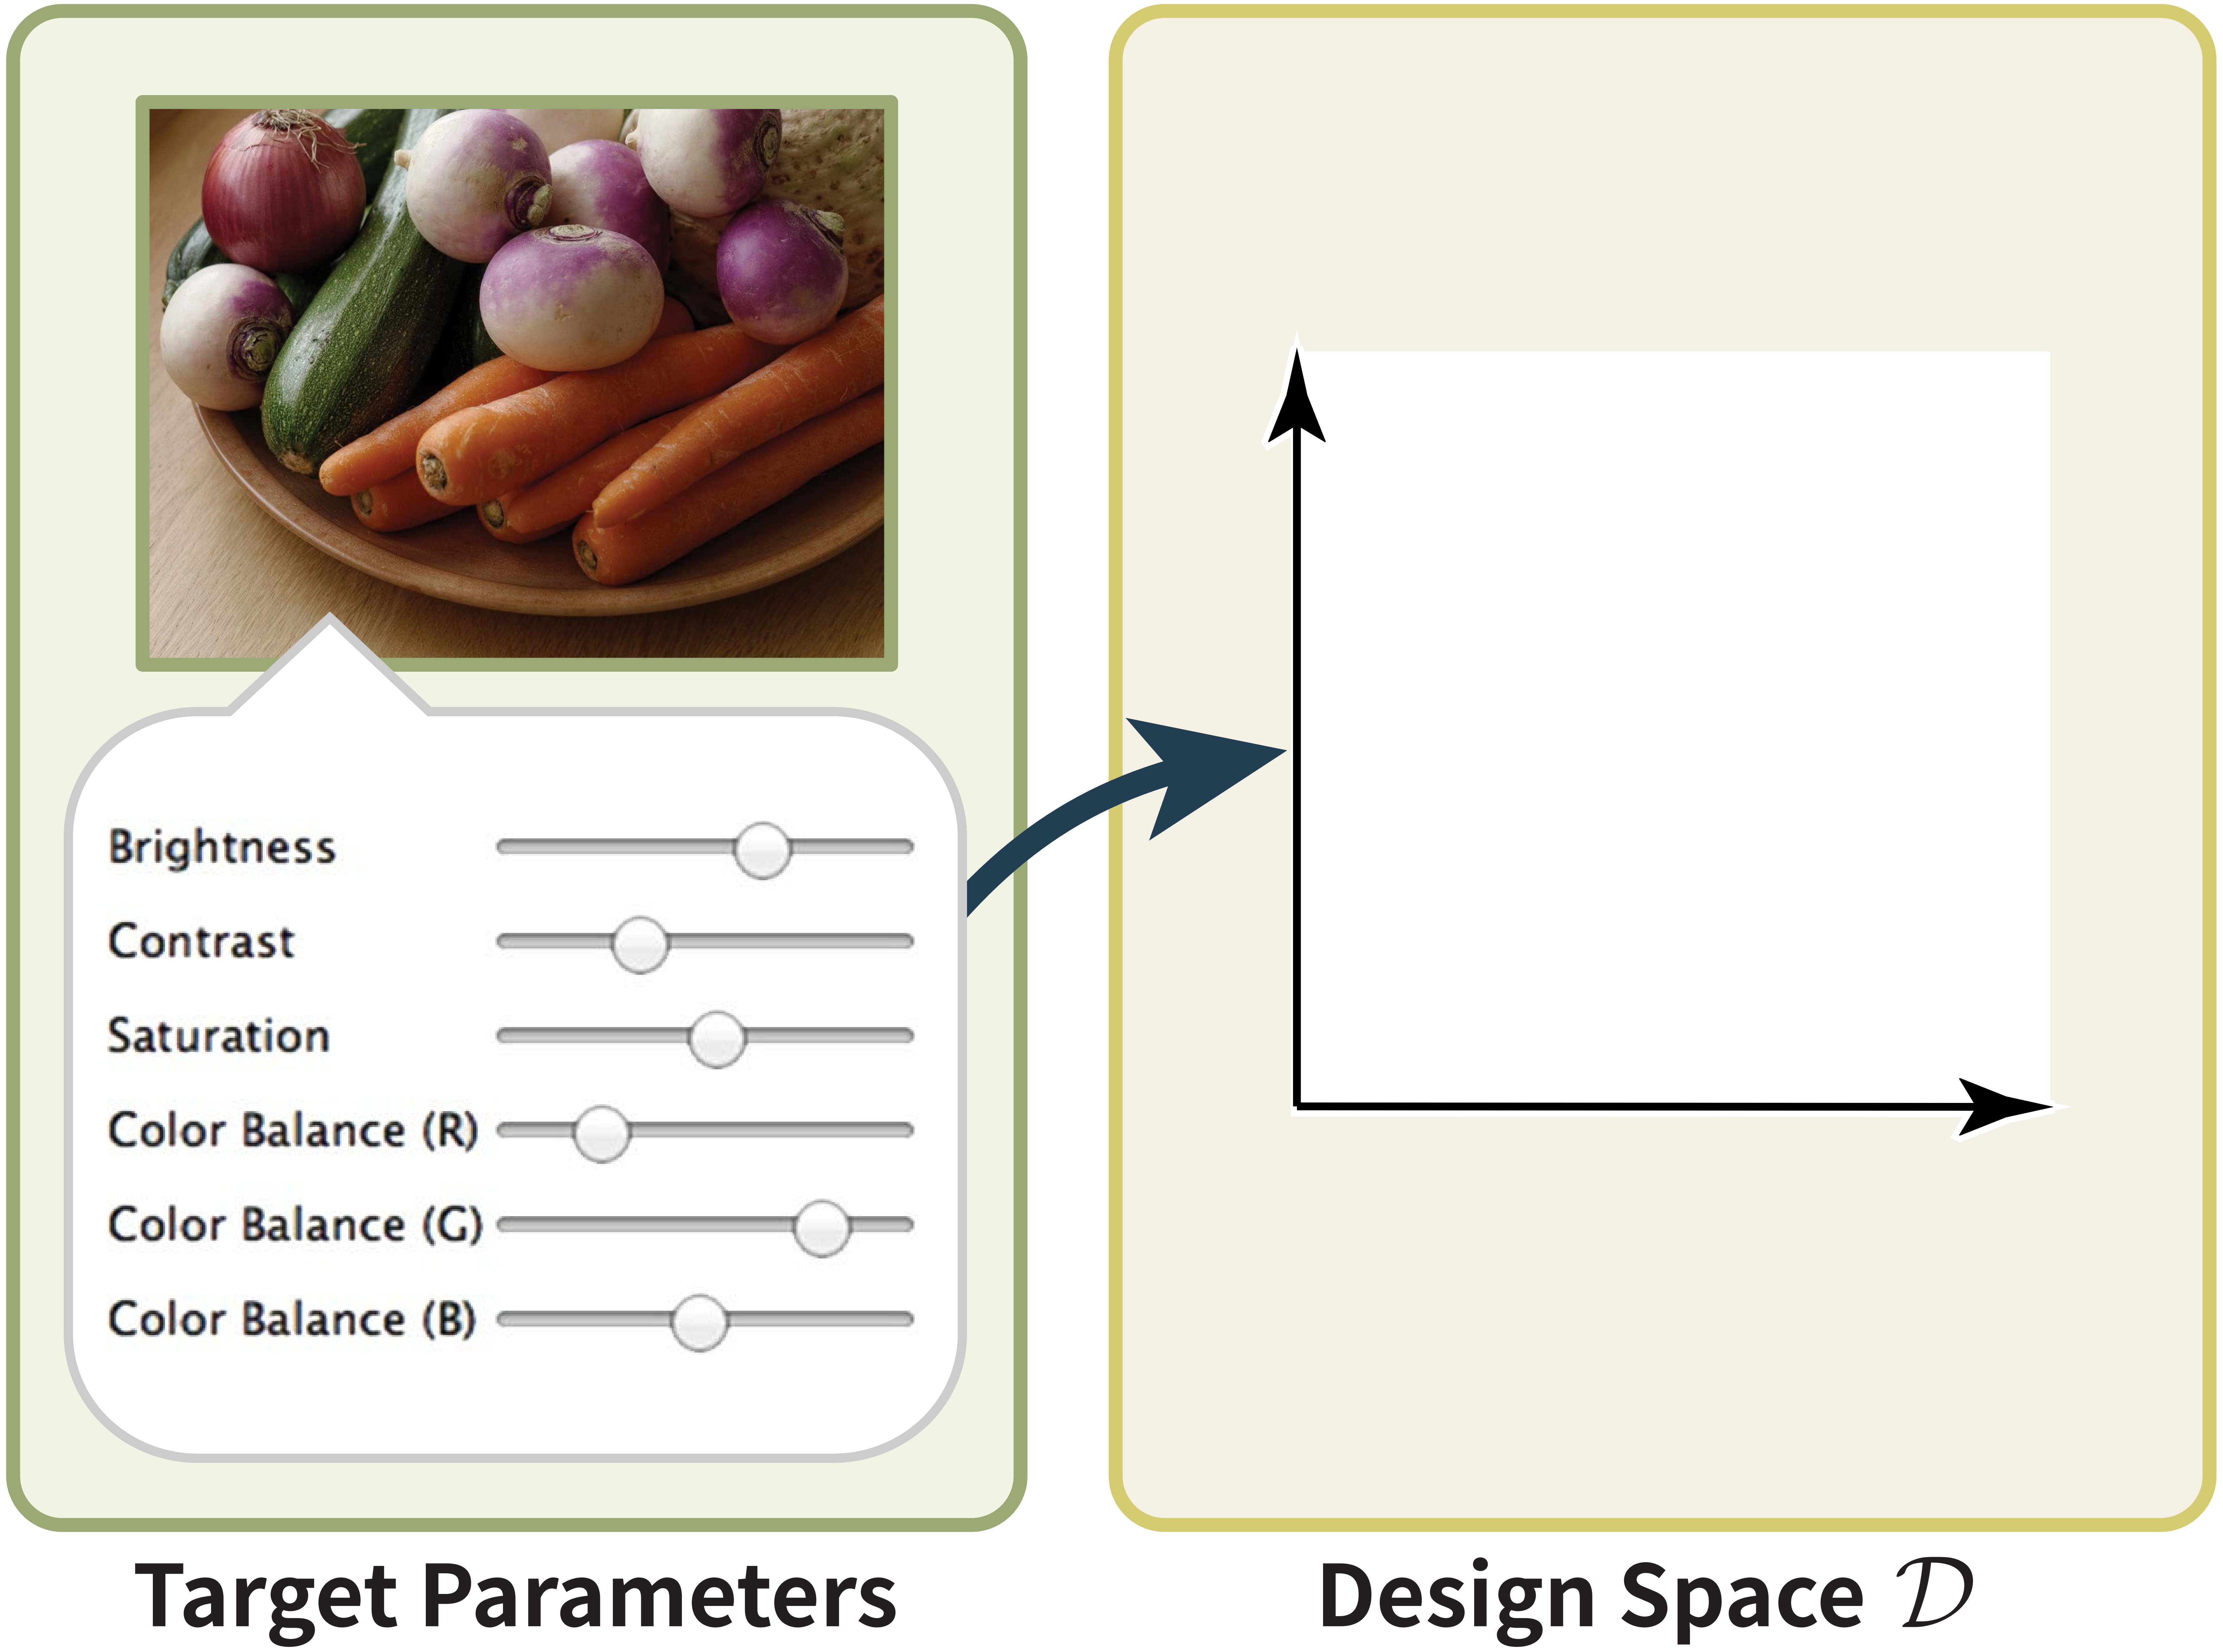
Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Target Parameters
Design Space $\mathcal{D}$

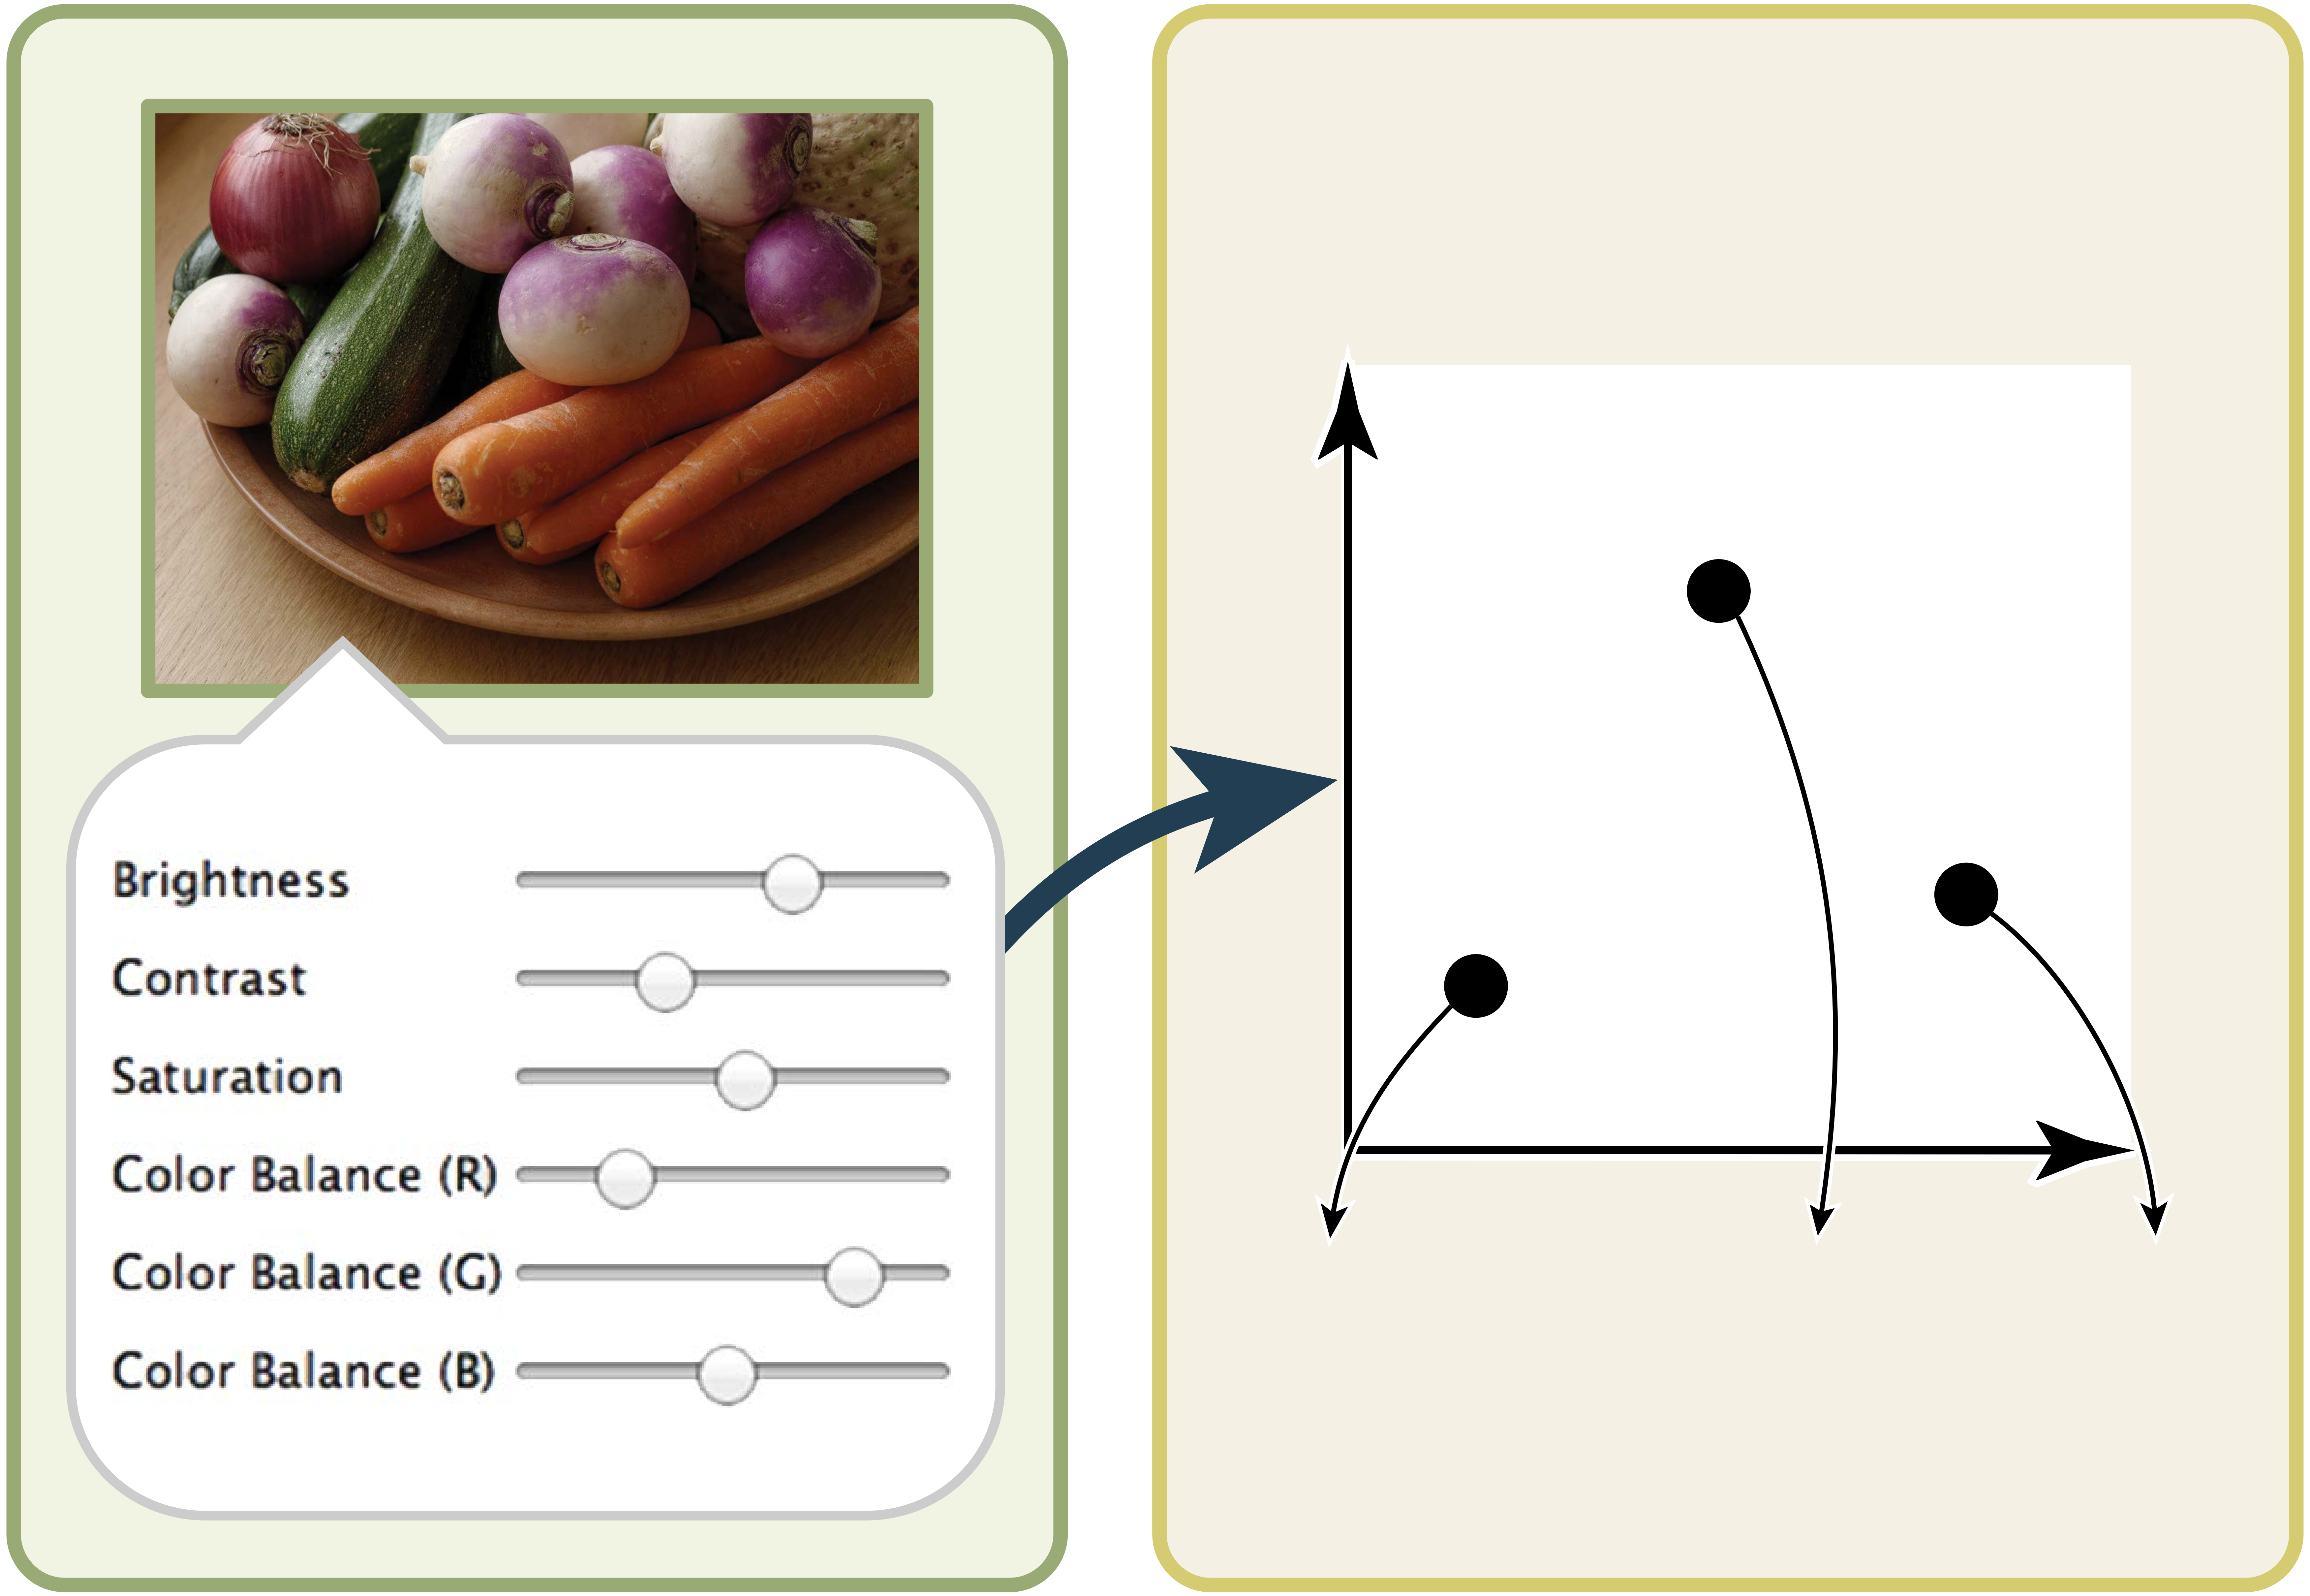

Target Parameters **Design Space $\mathcal{D}$**

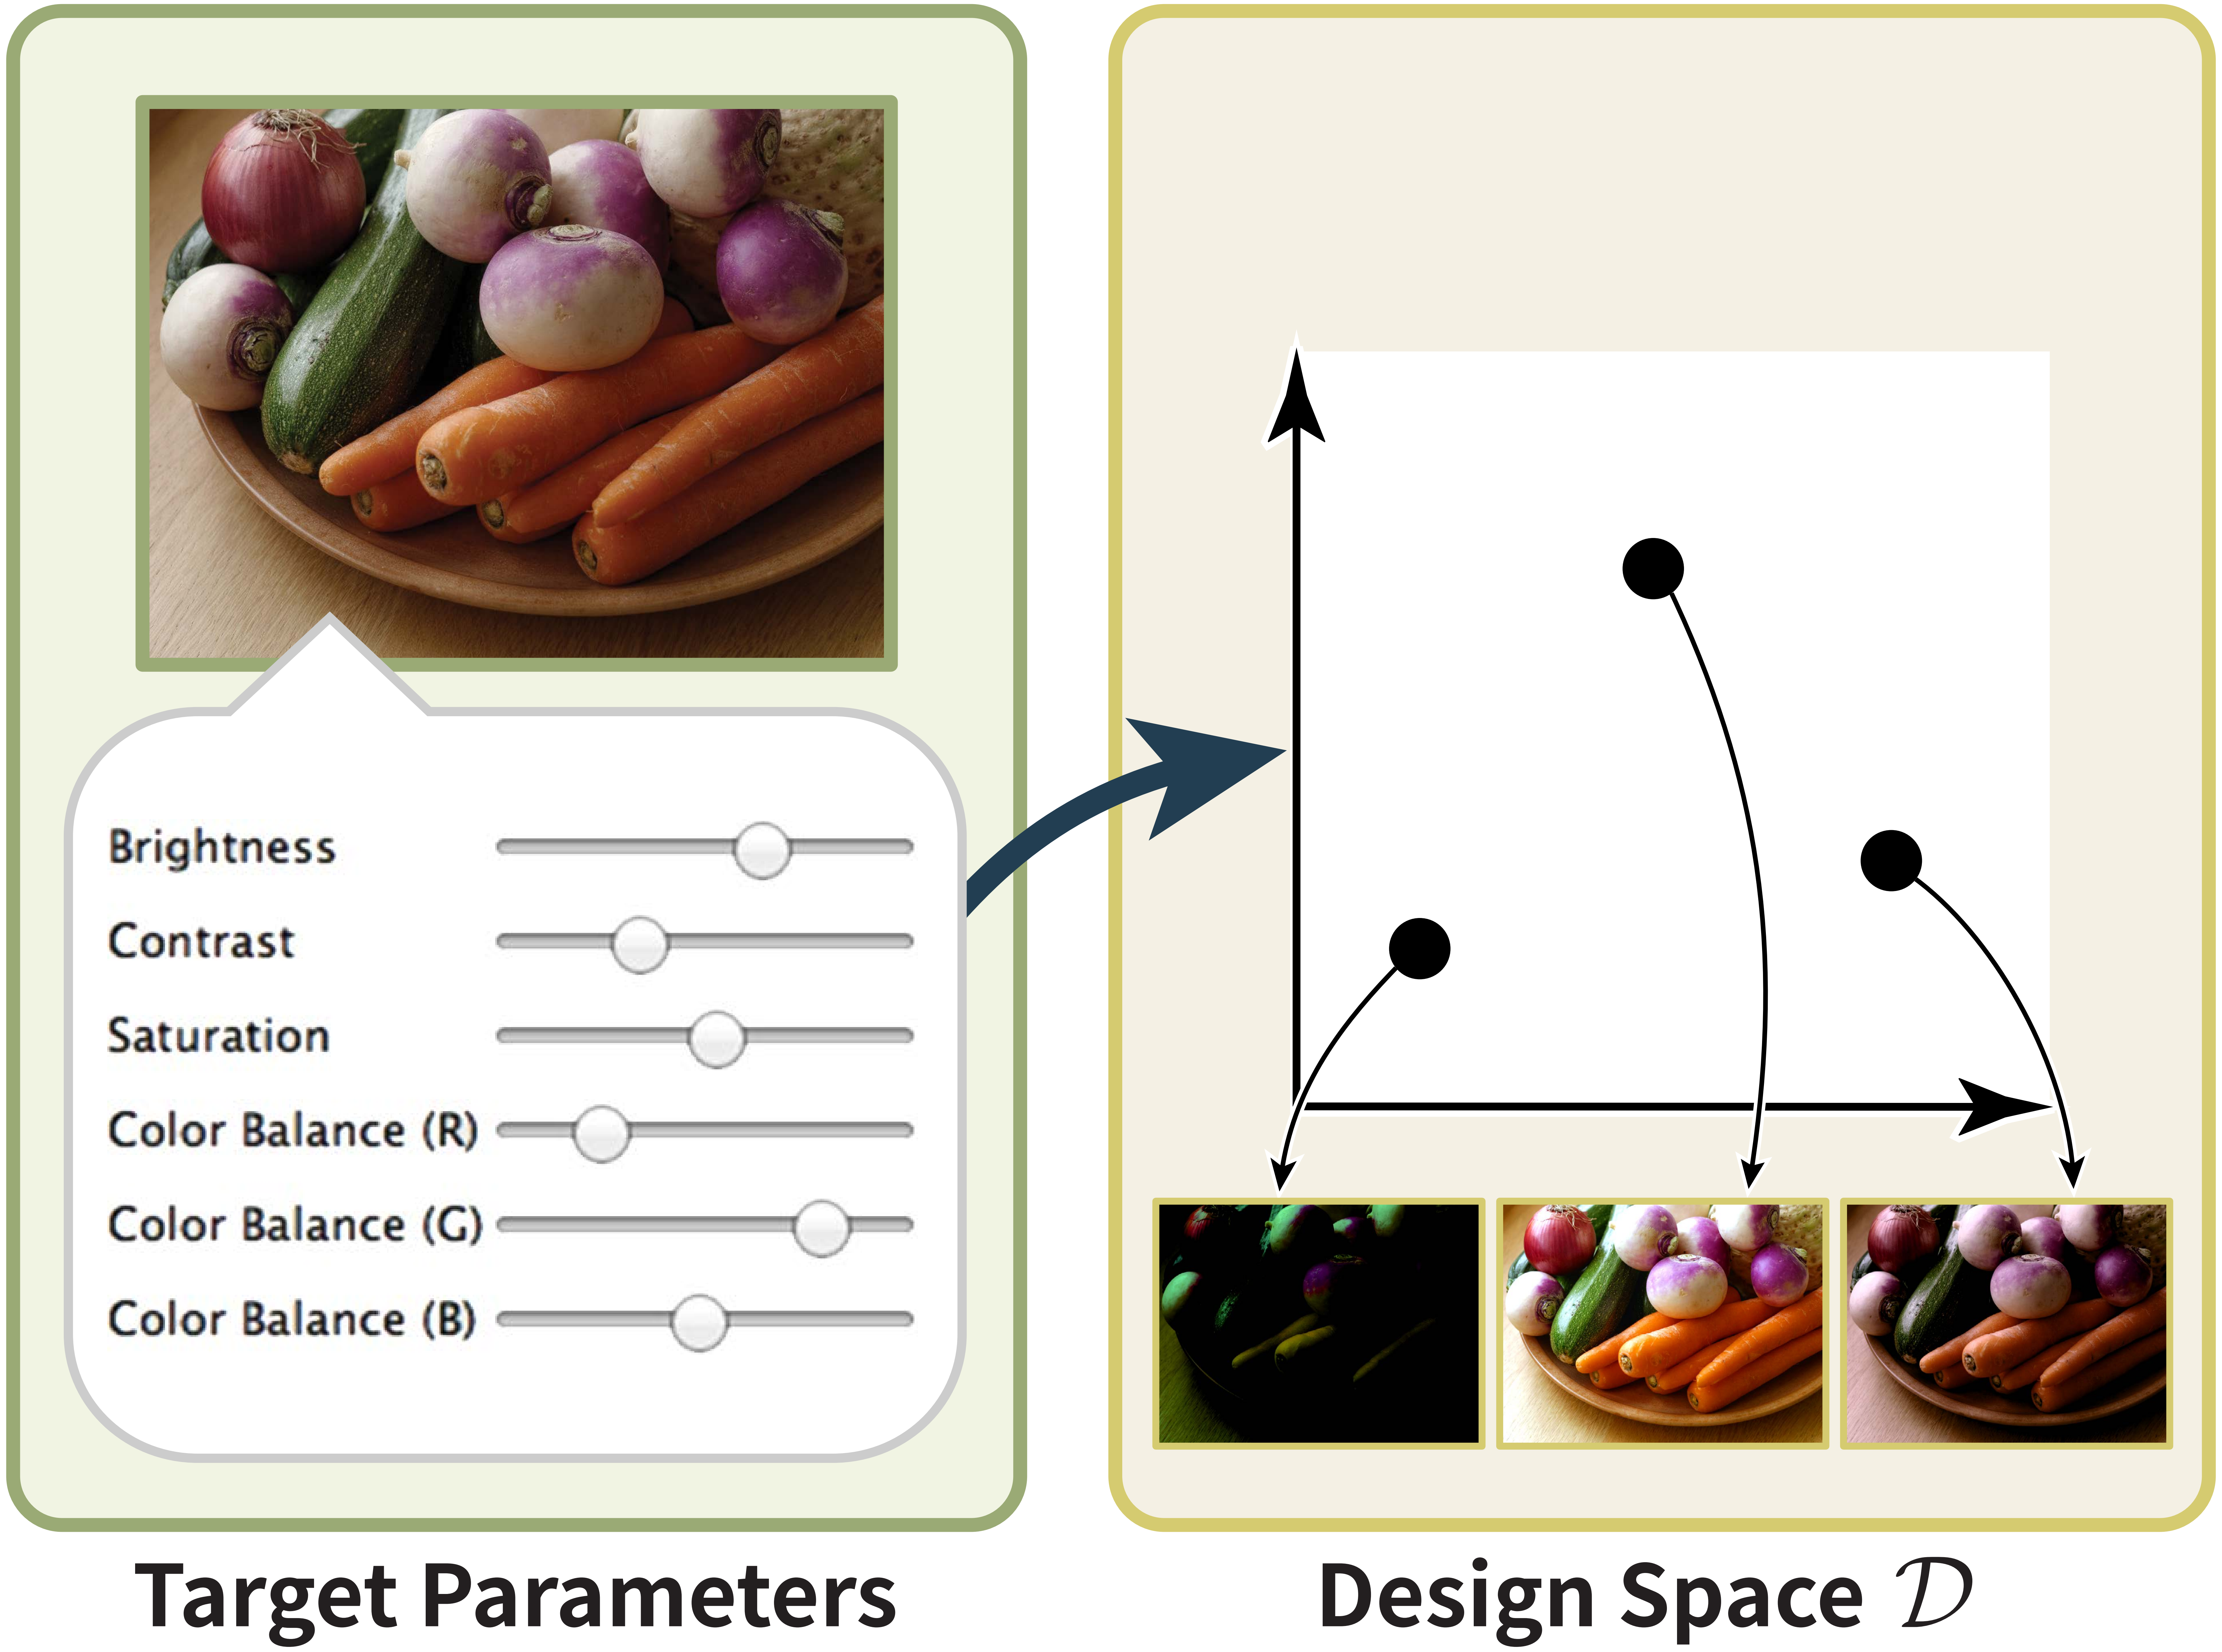
Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Target Parameters
Design Space $\mathcal{D}$

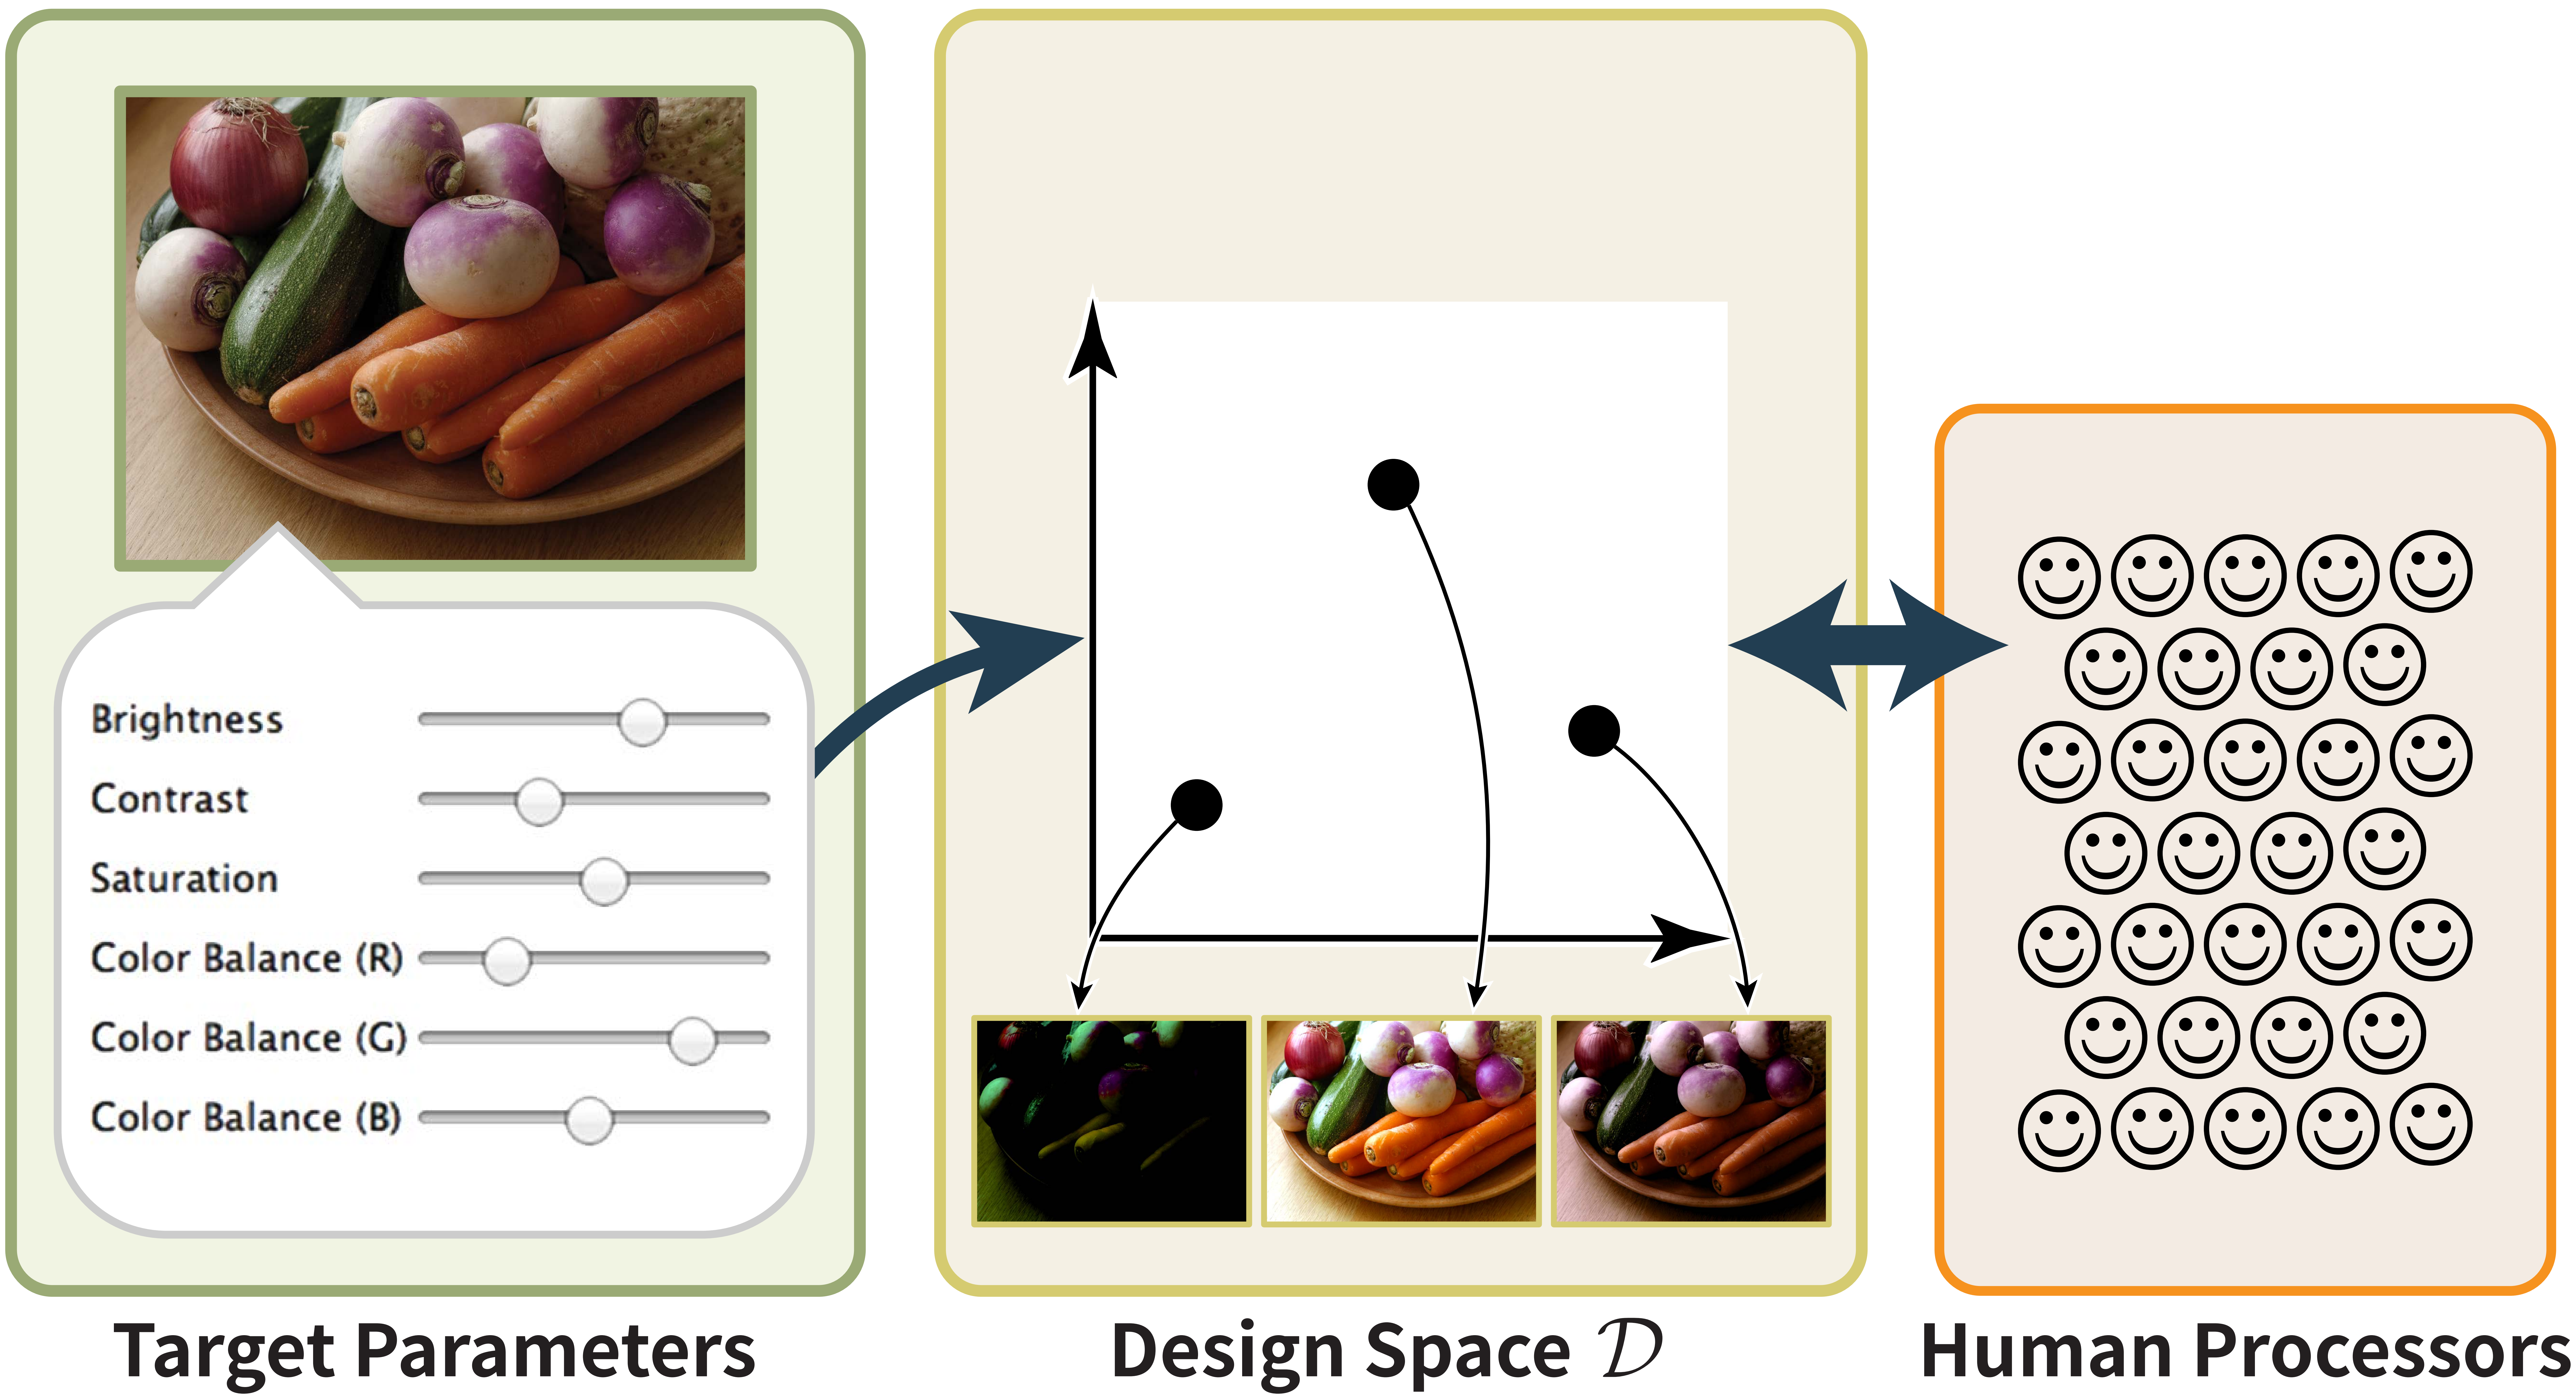
Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Target Parameters
Design Space $\mathcal{D}$
Human Processors

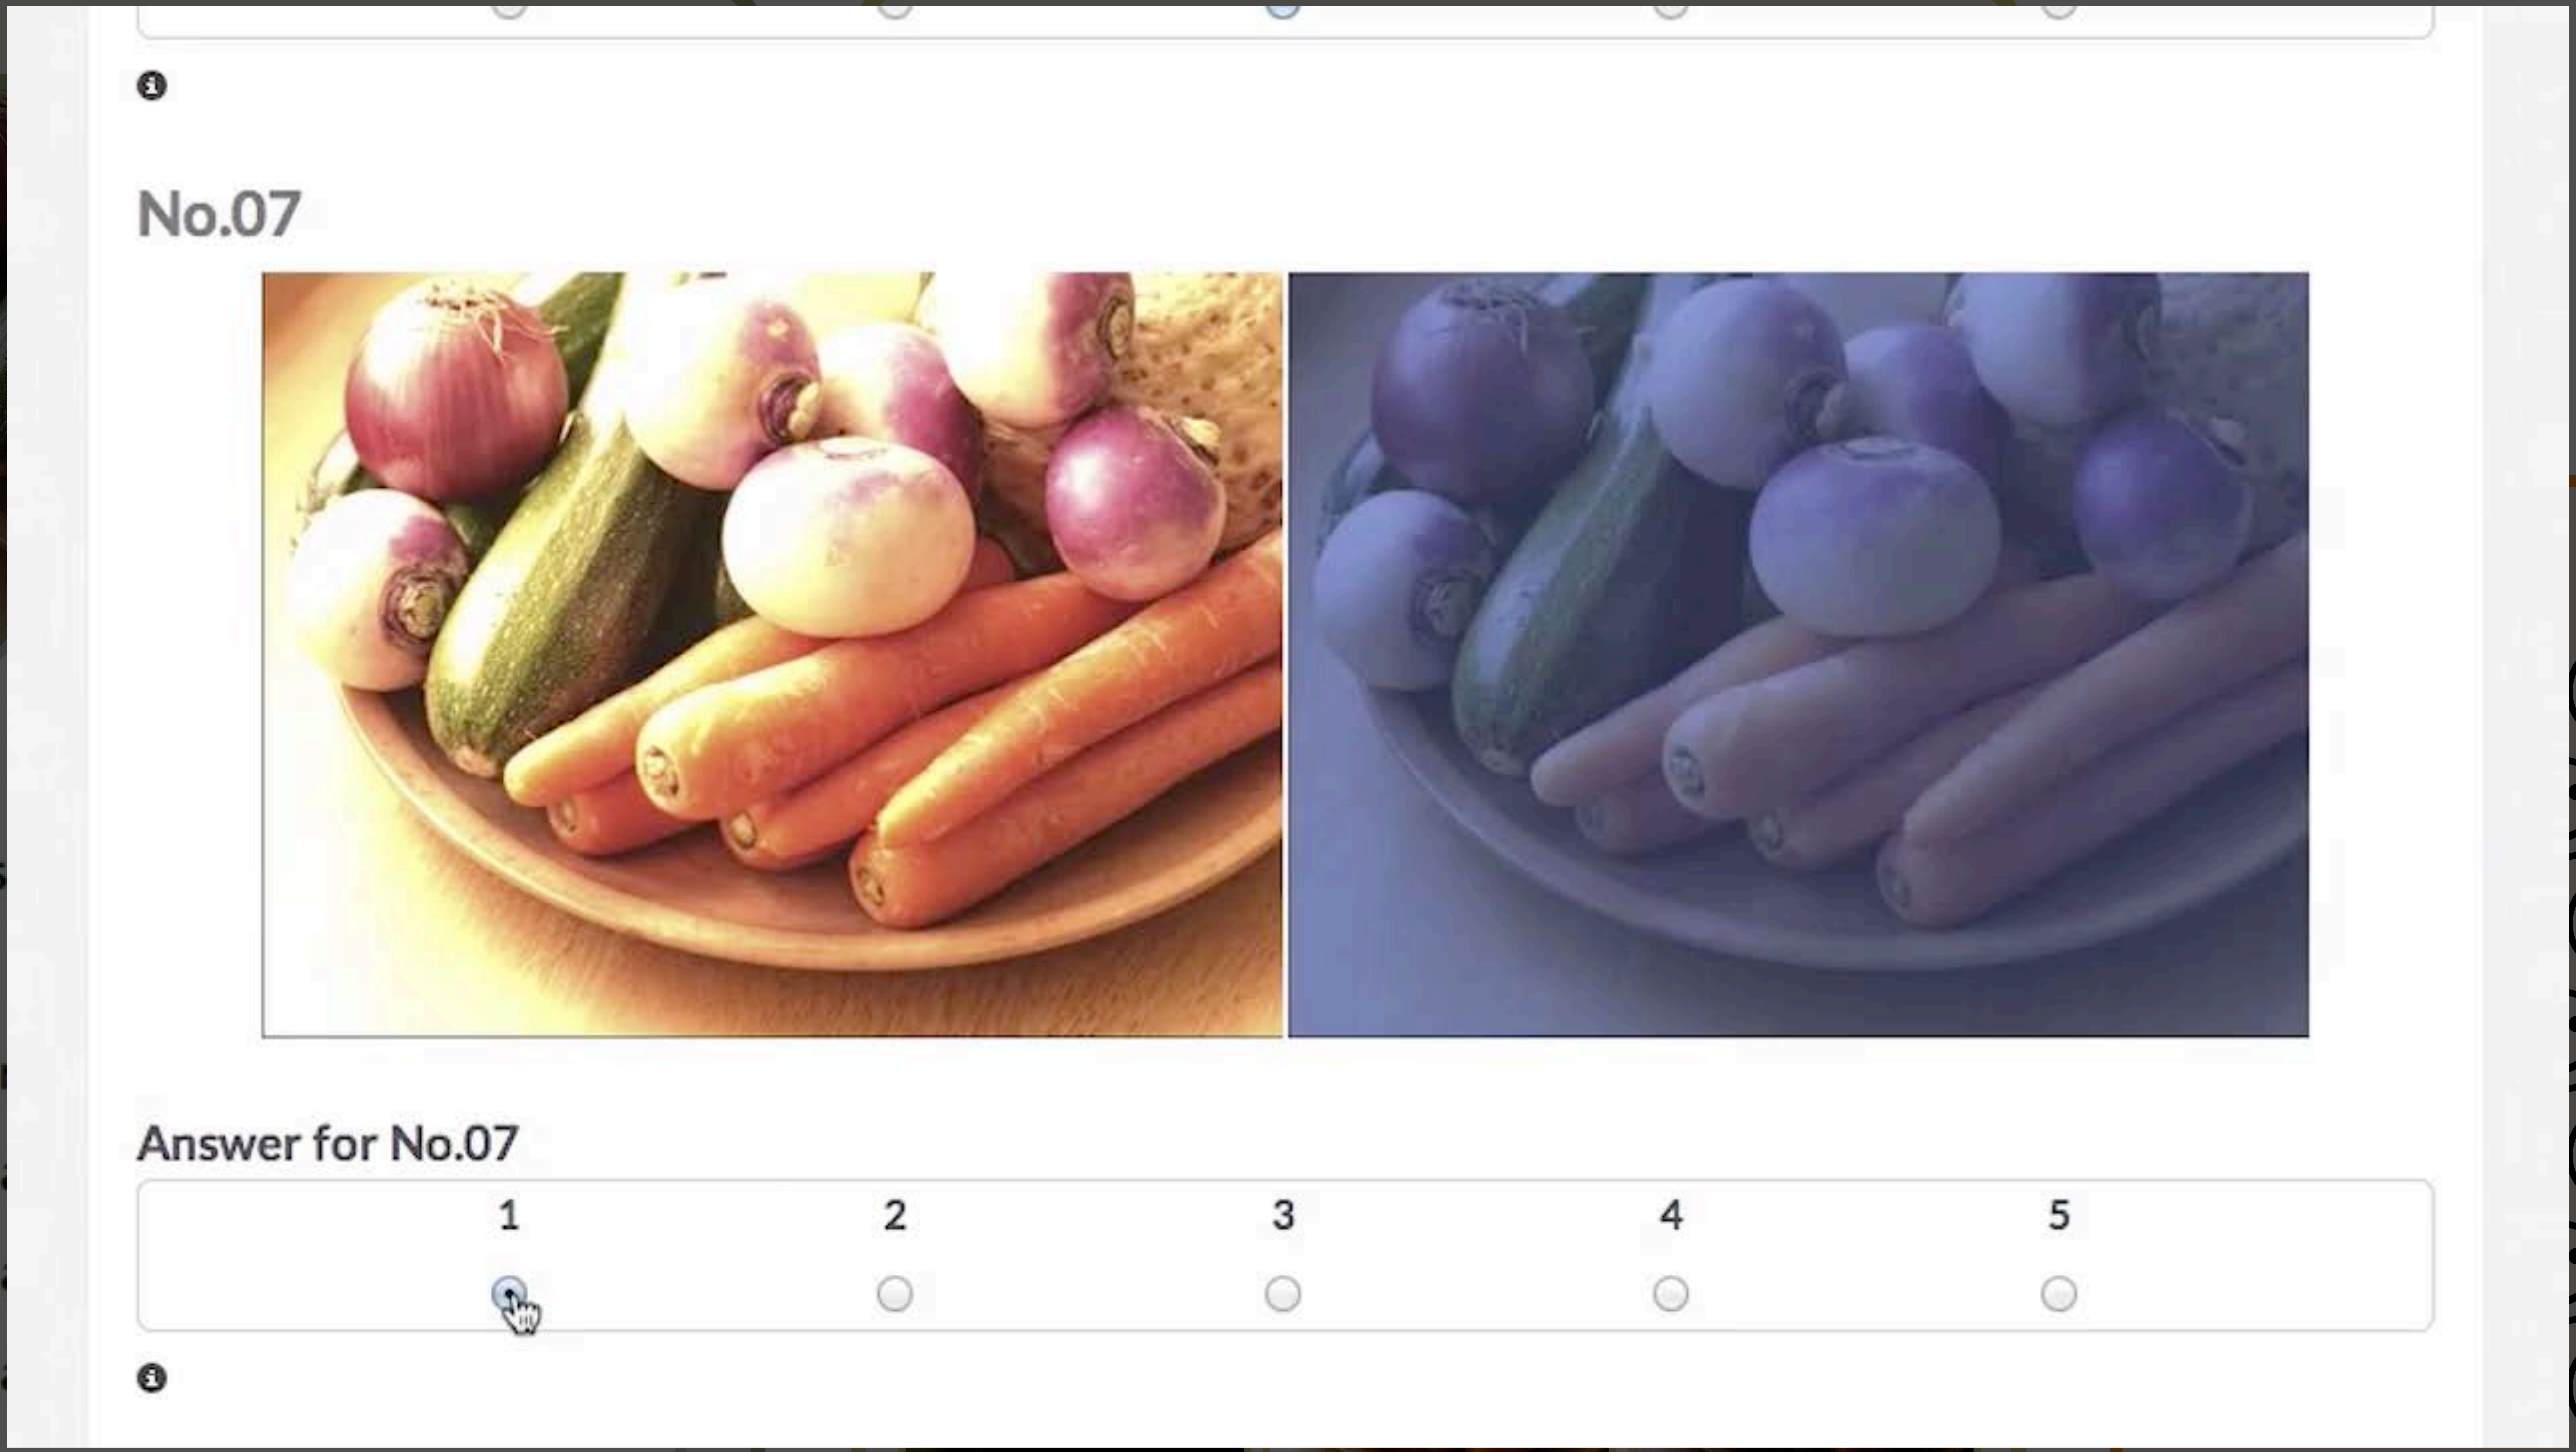

Preference data generation [pairwise comparison]

(c.f., 2000 comparisons, 4 USD, 30 min)

Target Parameters

Human Processors

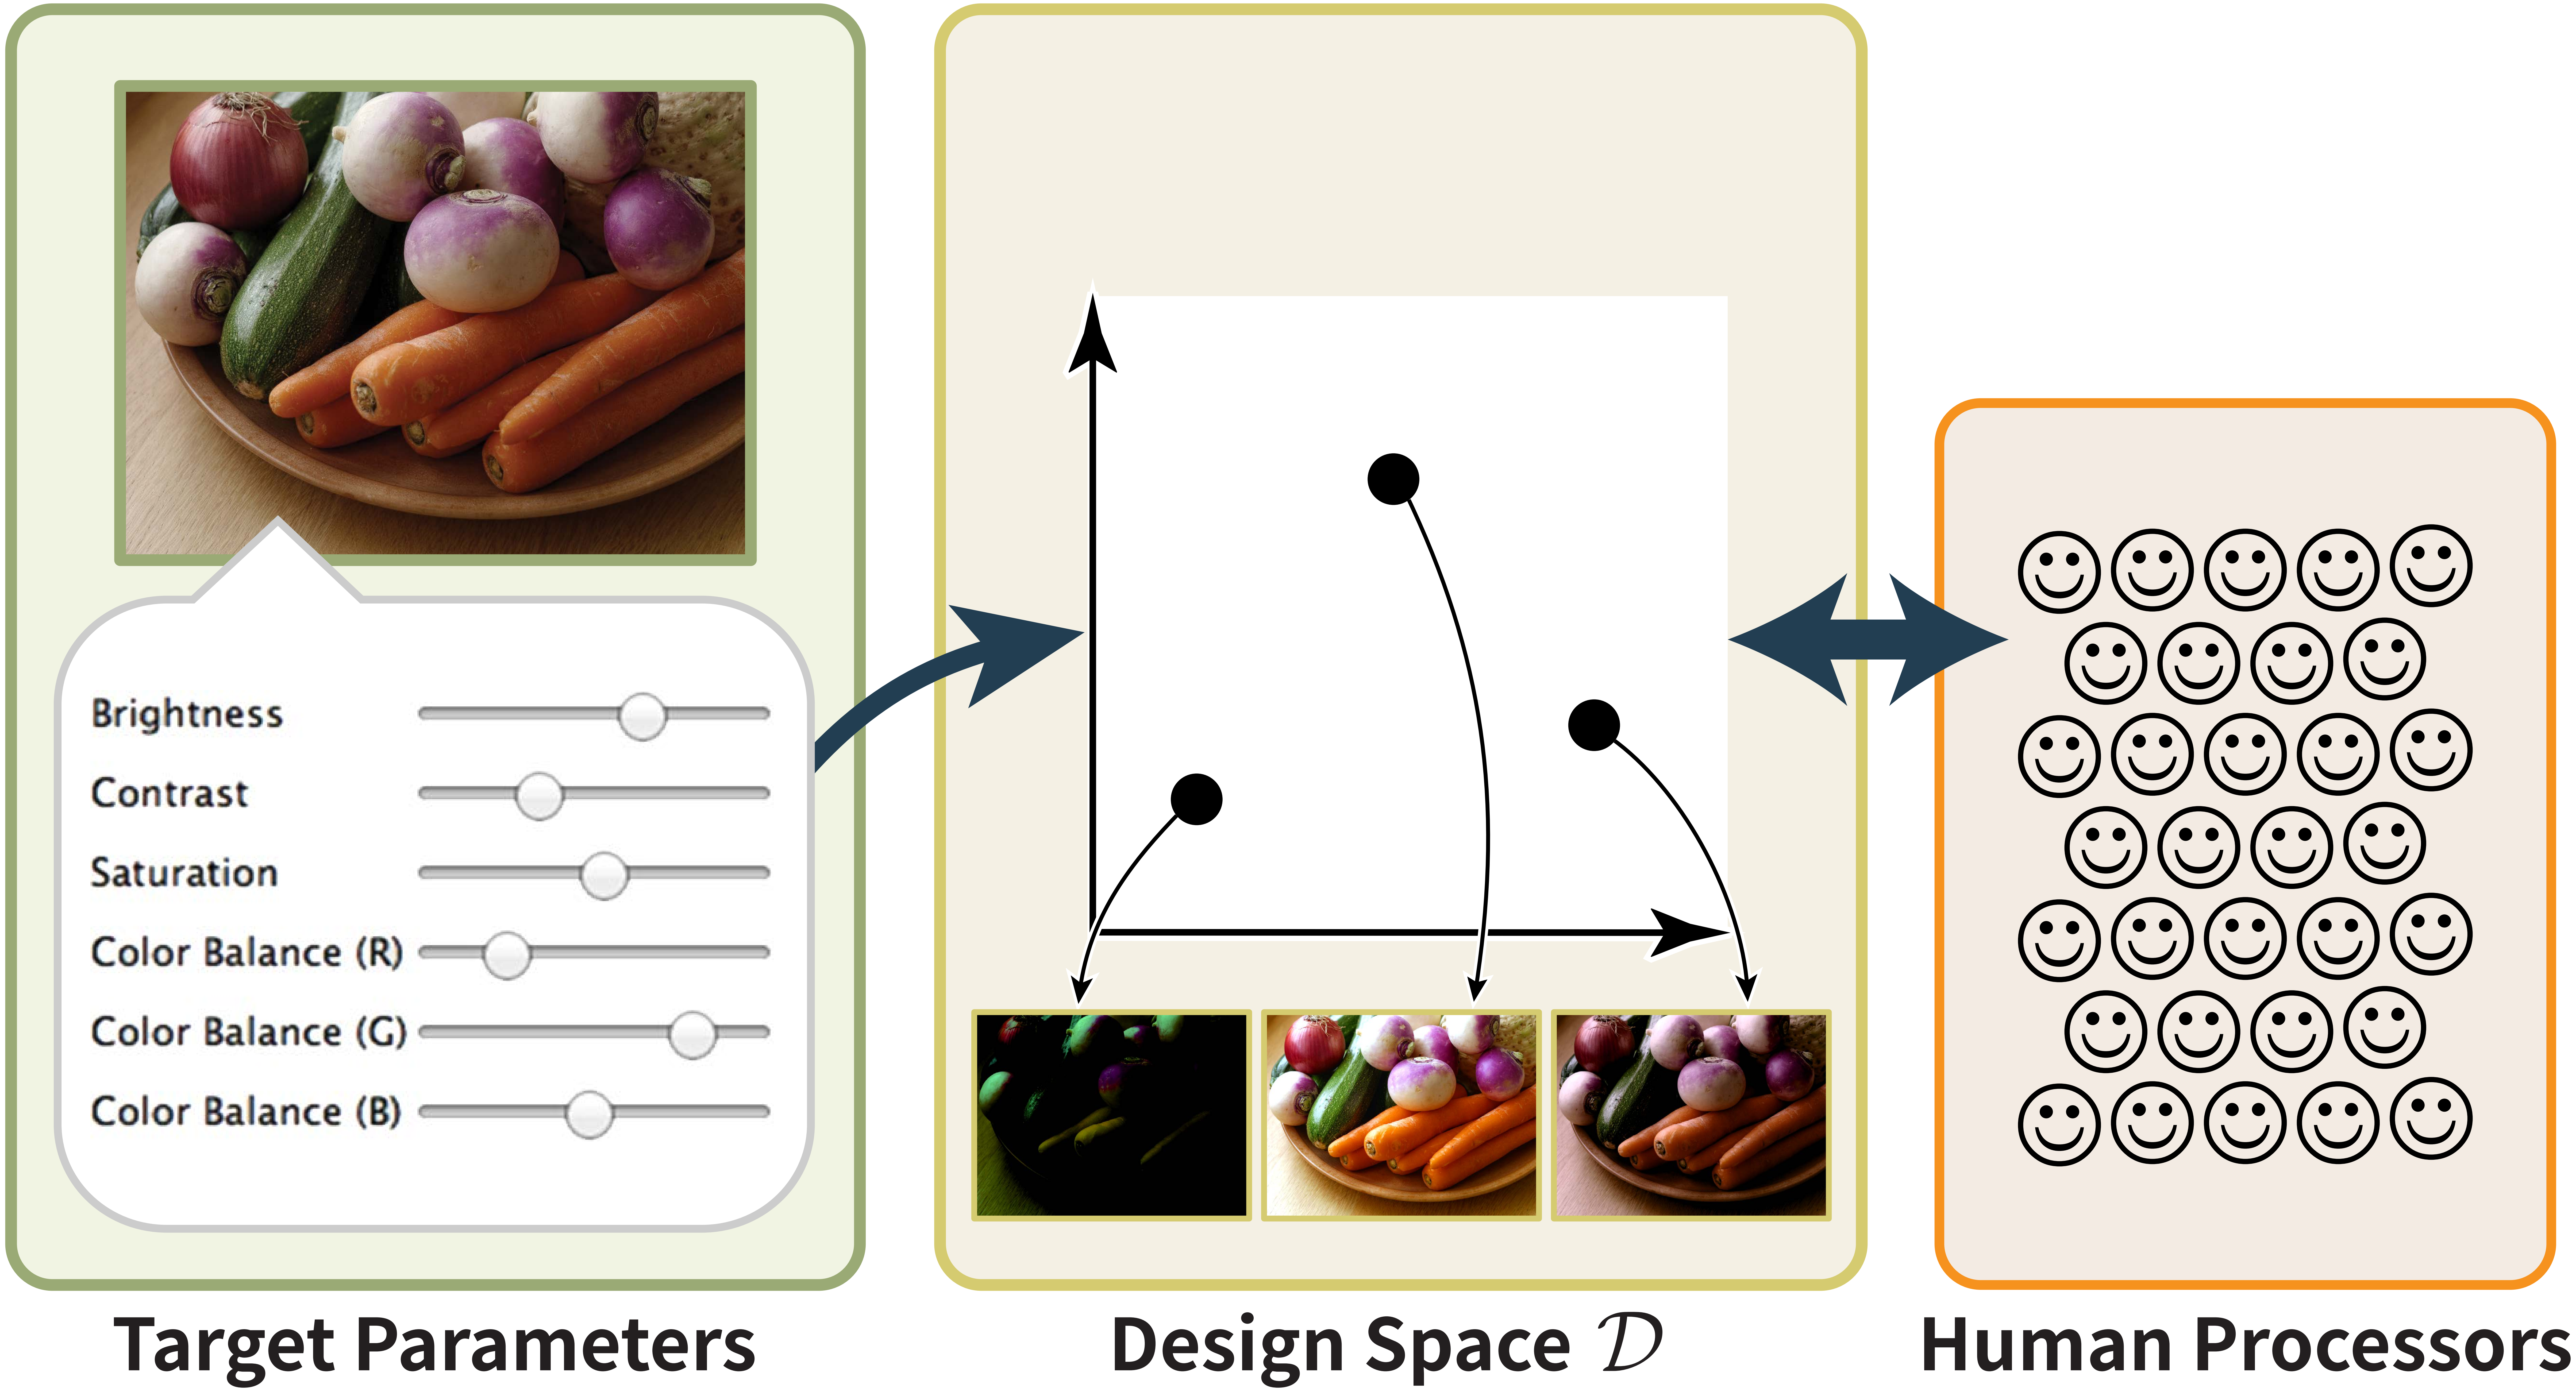
Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Target Parameters
Design Space $\mathcal{D}$
Human Processors

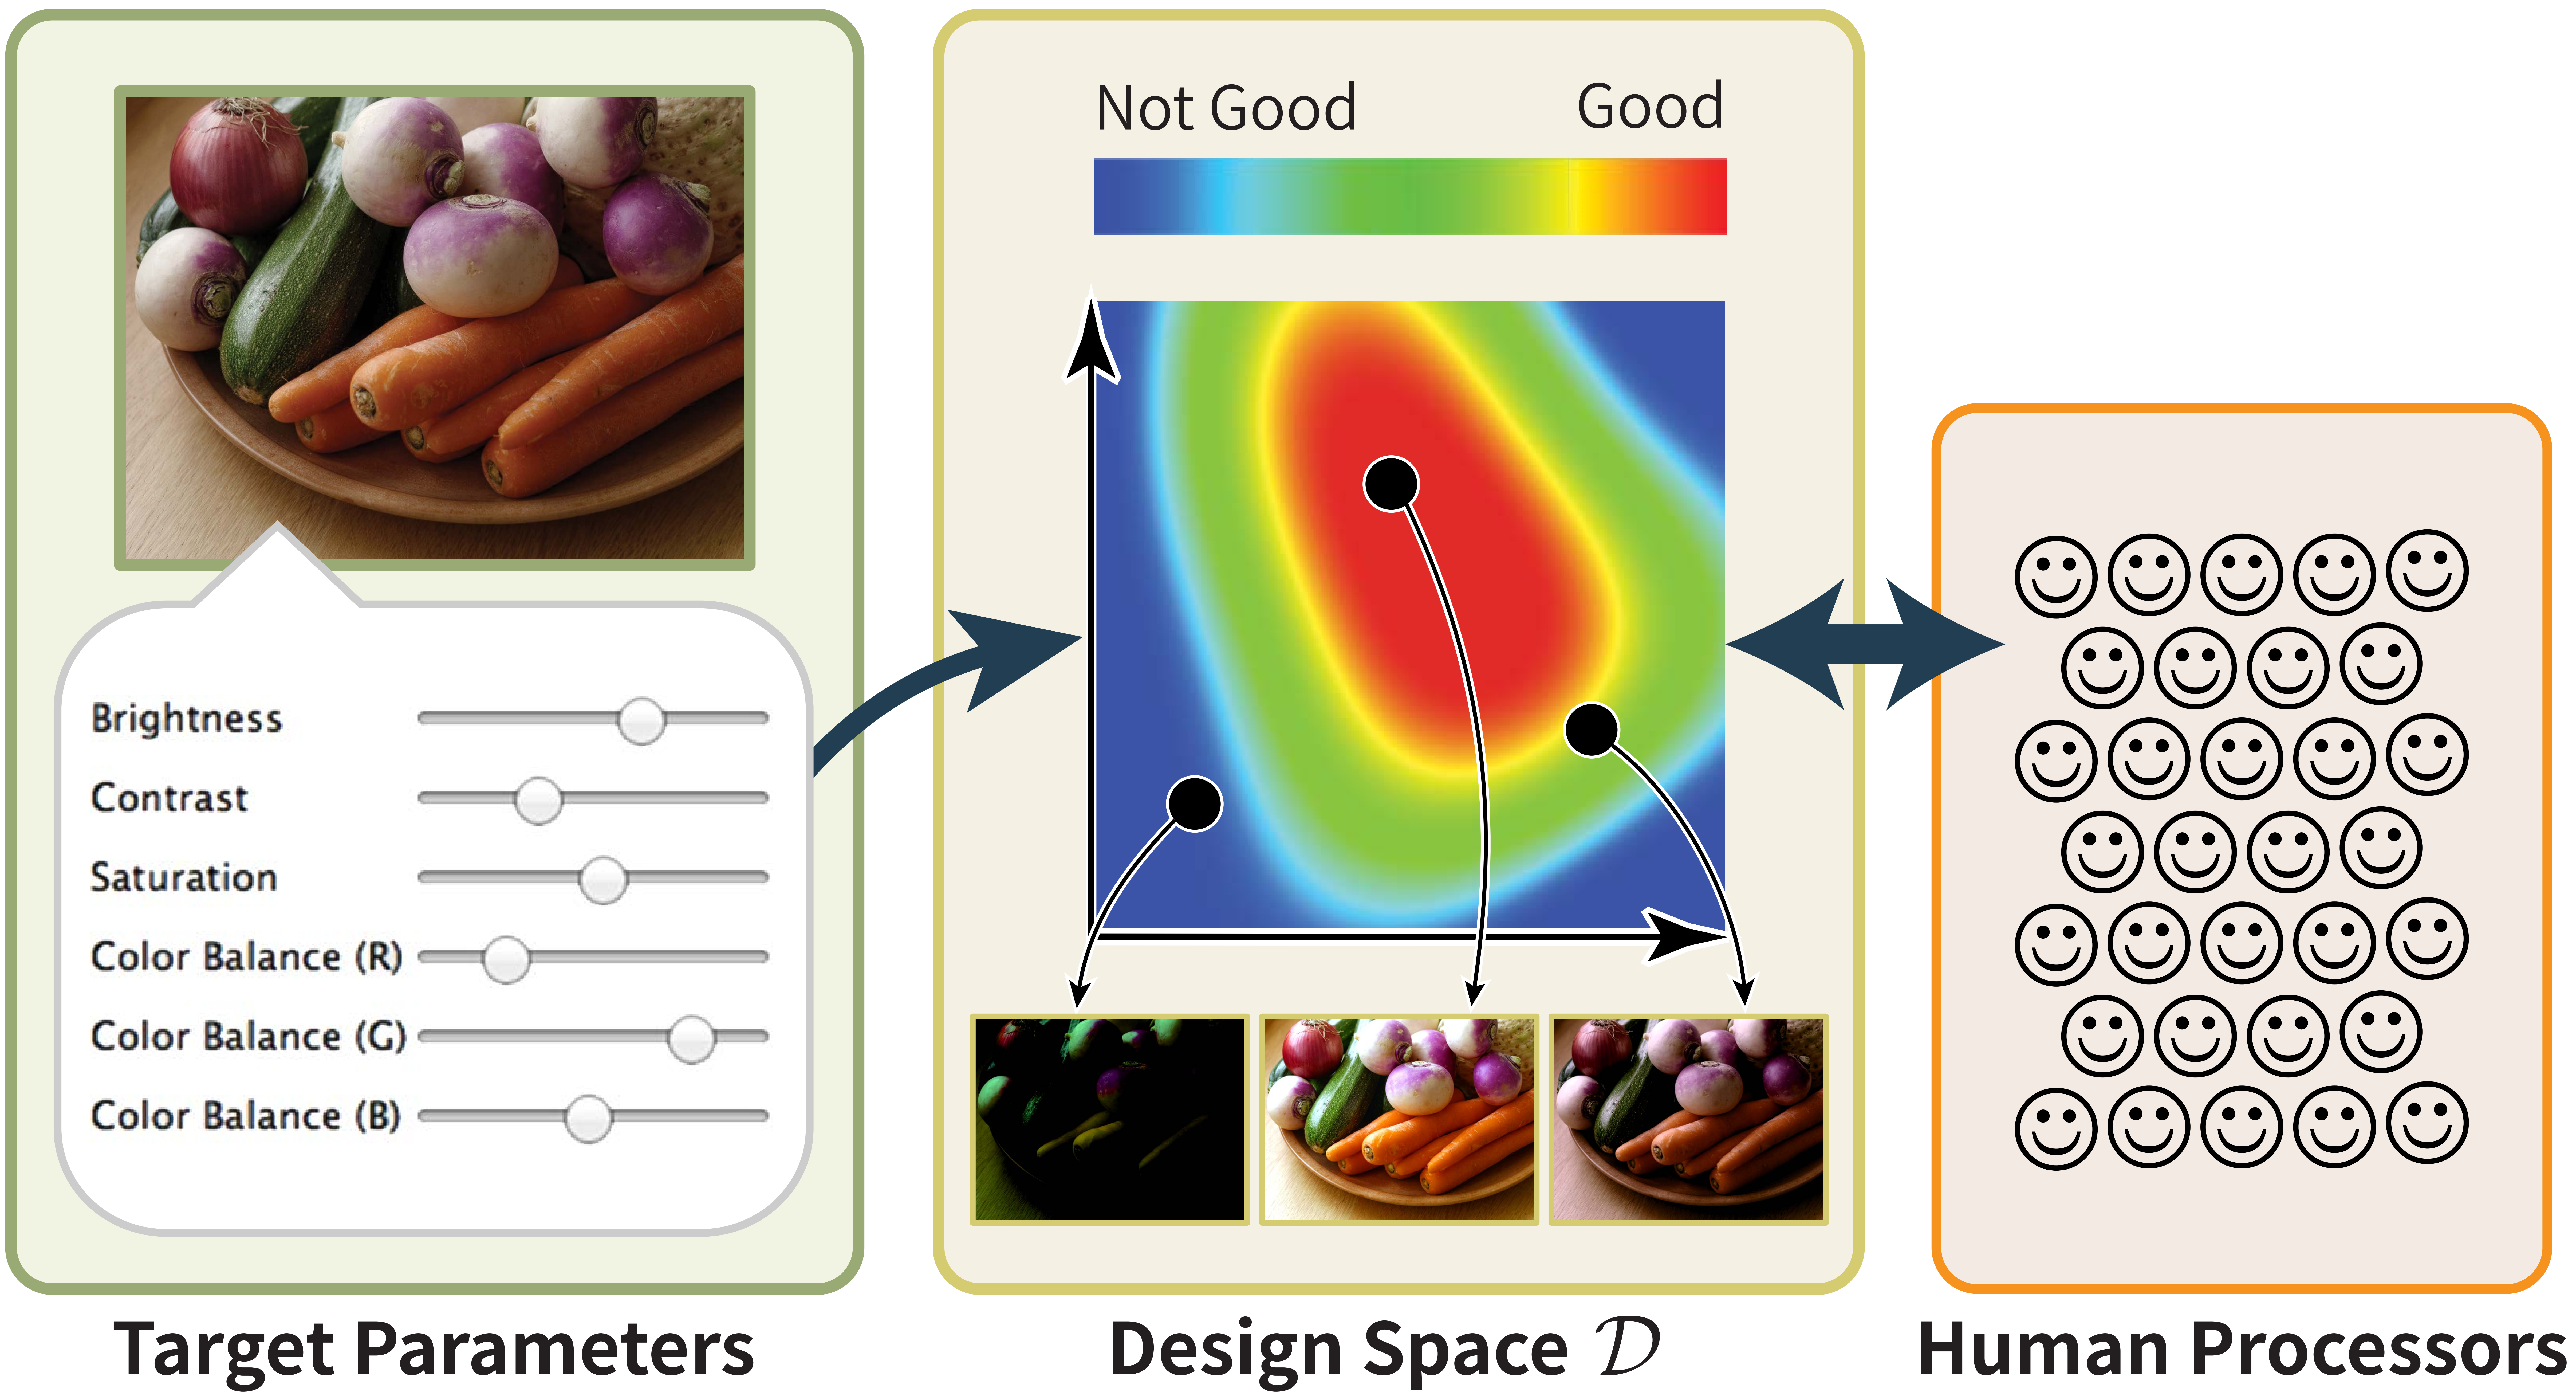
Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Not Good
Good
Target Parameters
Design Space $\mathcal{D}$
Human Processors

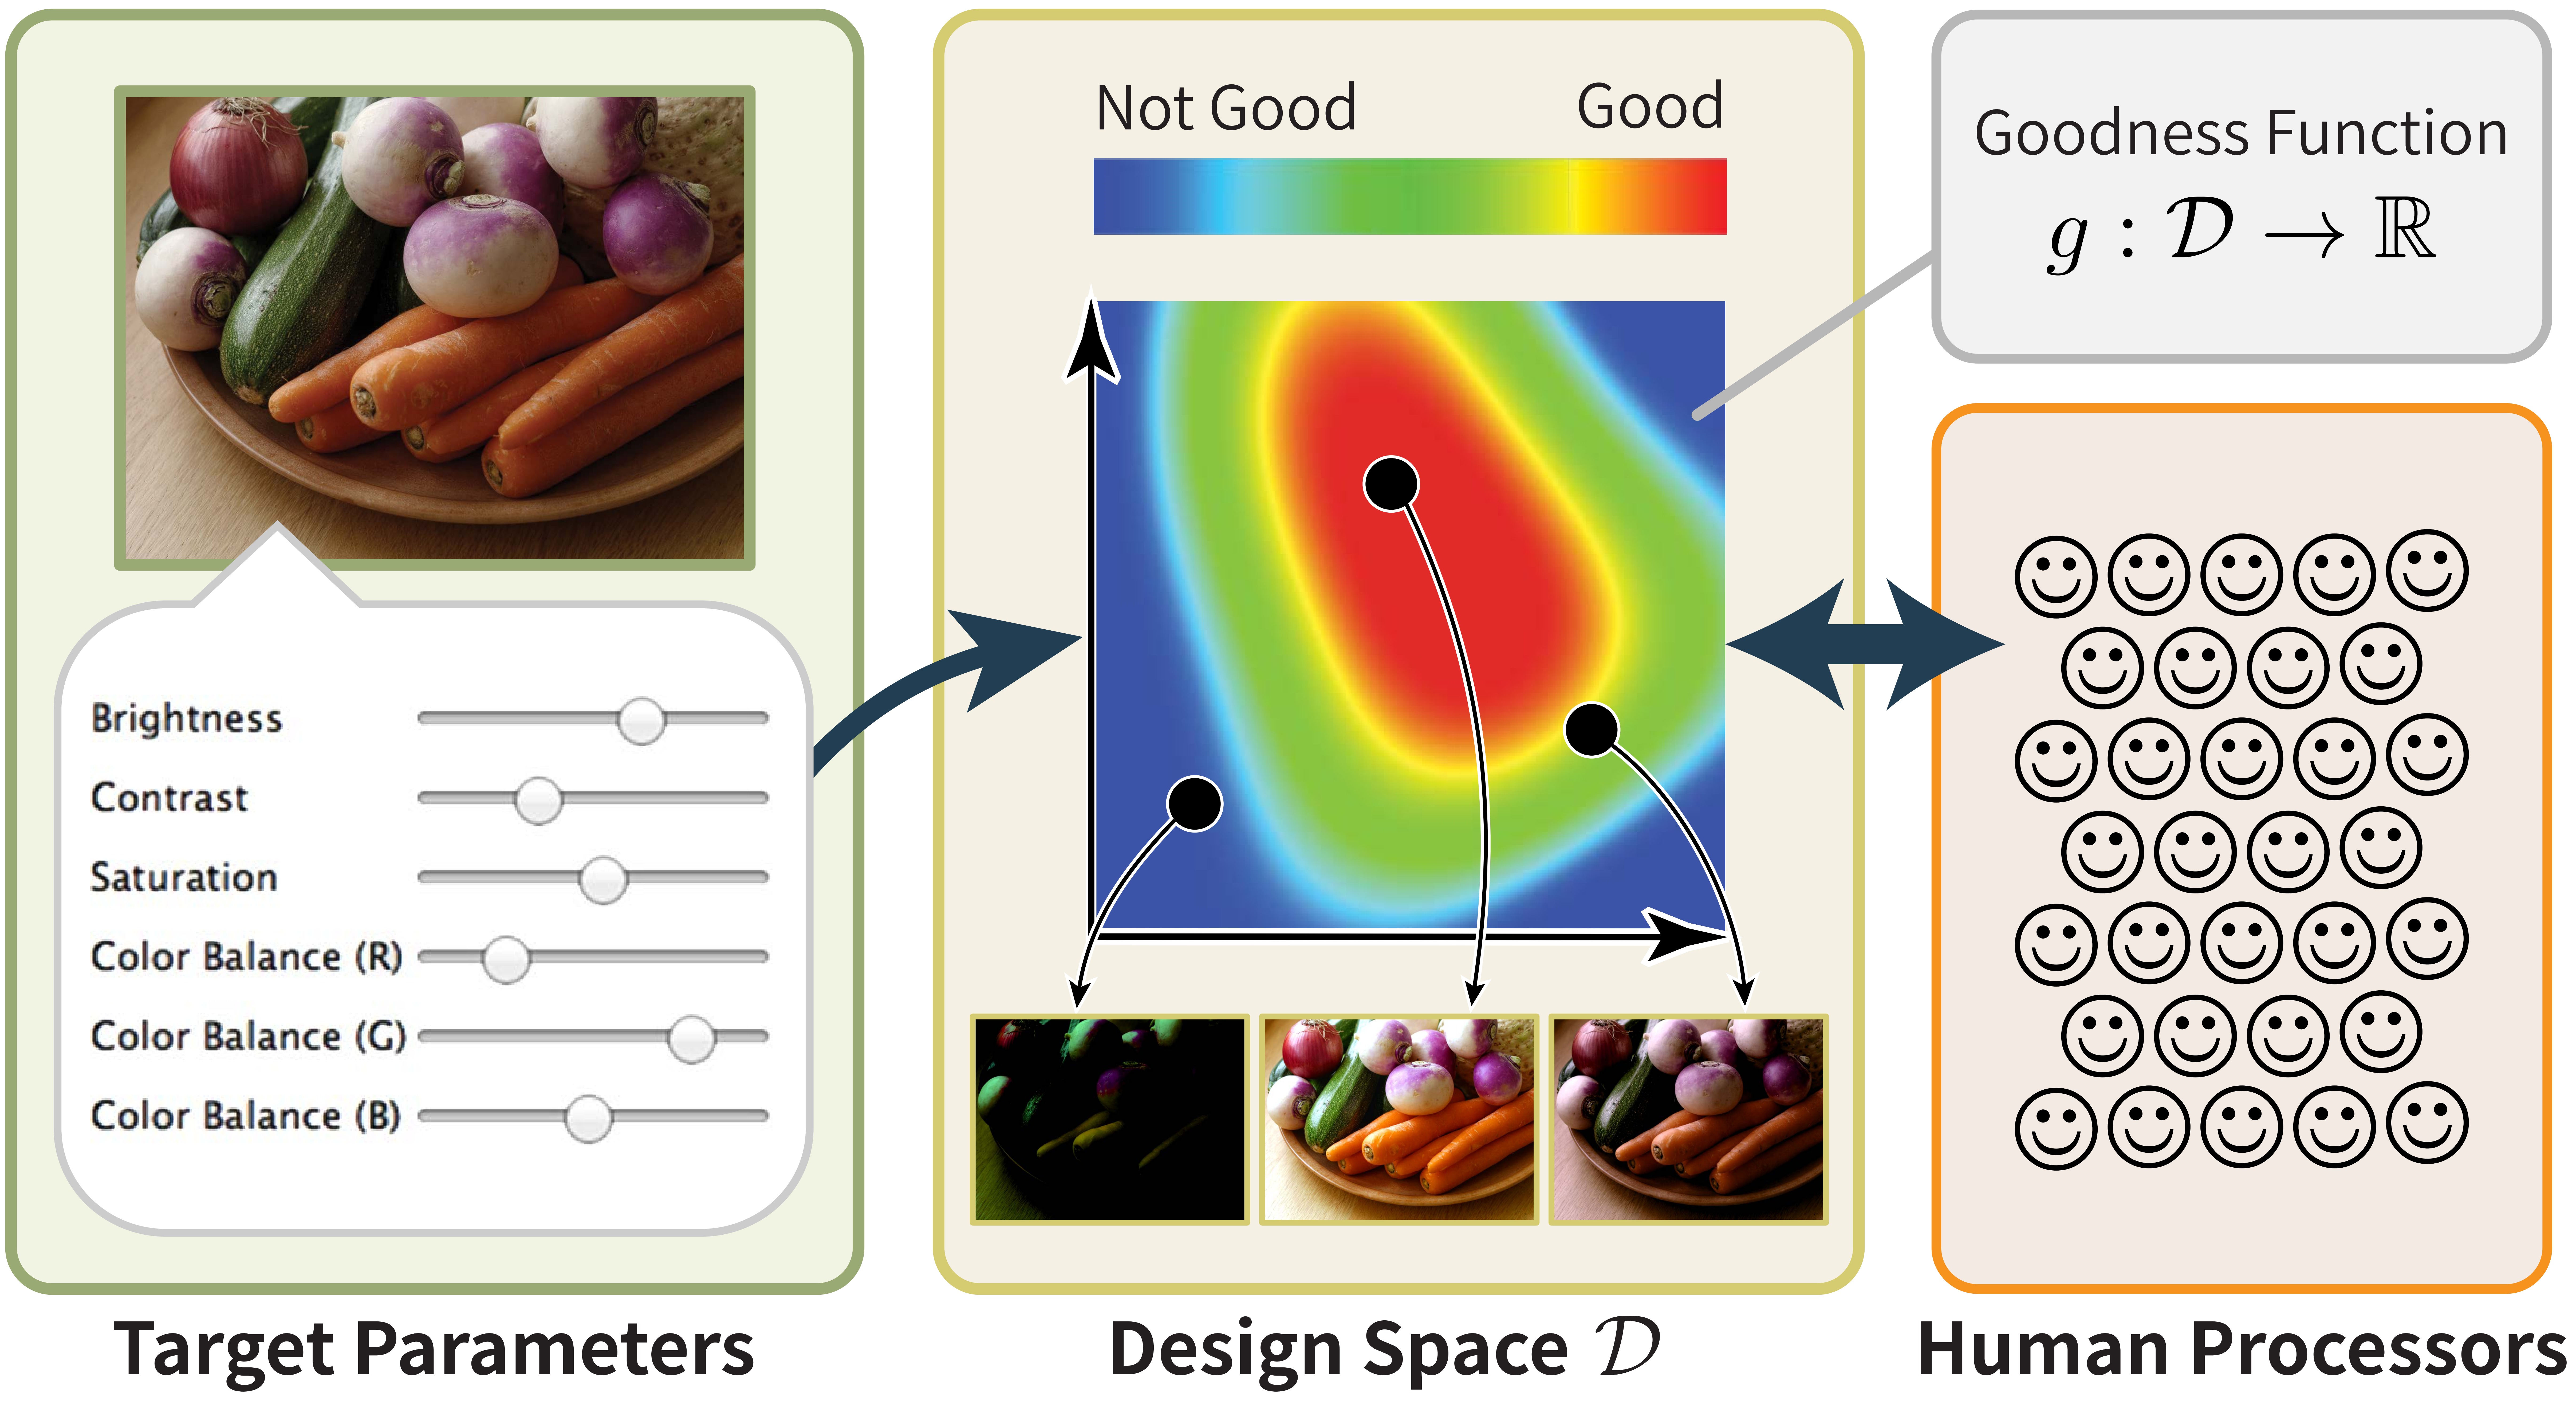

Brightness
Contrast
Saturation
Color Balance (R)
Color Balance (G)
Color Balance (B)
Not Good
Good
Goodness Function
$g : \mathcal{D} \to \mathbb{R}$
Target Parameters
Design Space $\mathcal{D}$
Human Processors

Intelligent Tools by Goodness Function Estimation

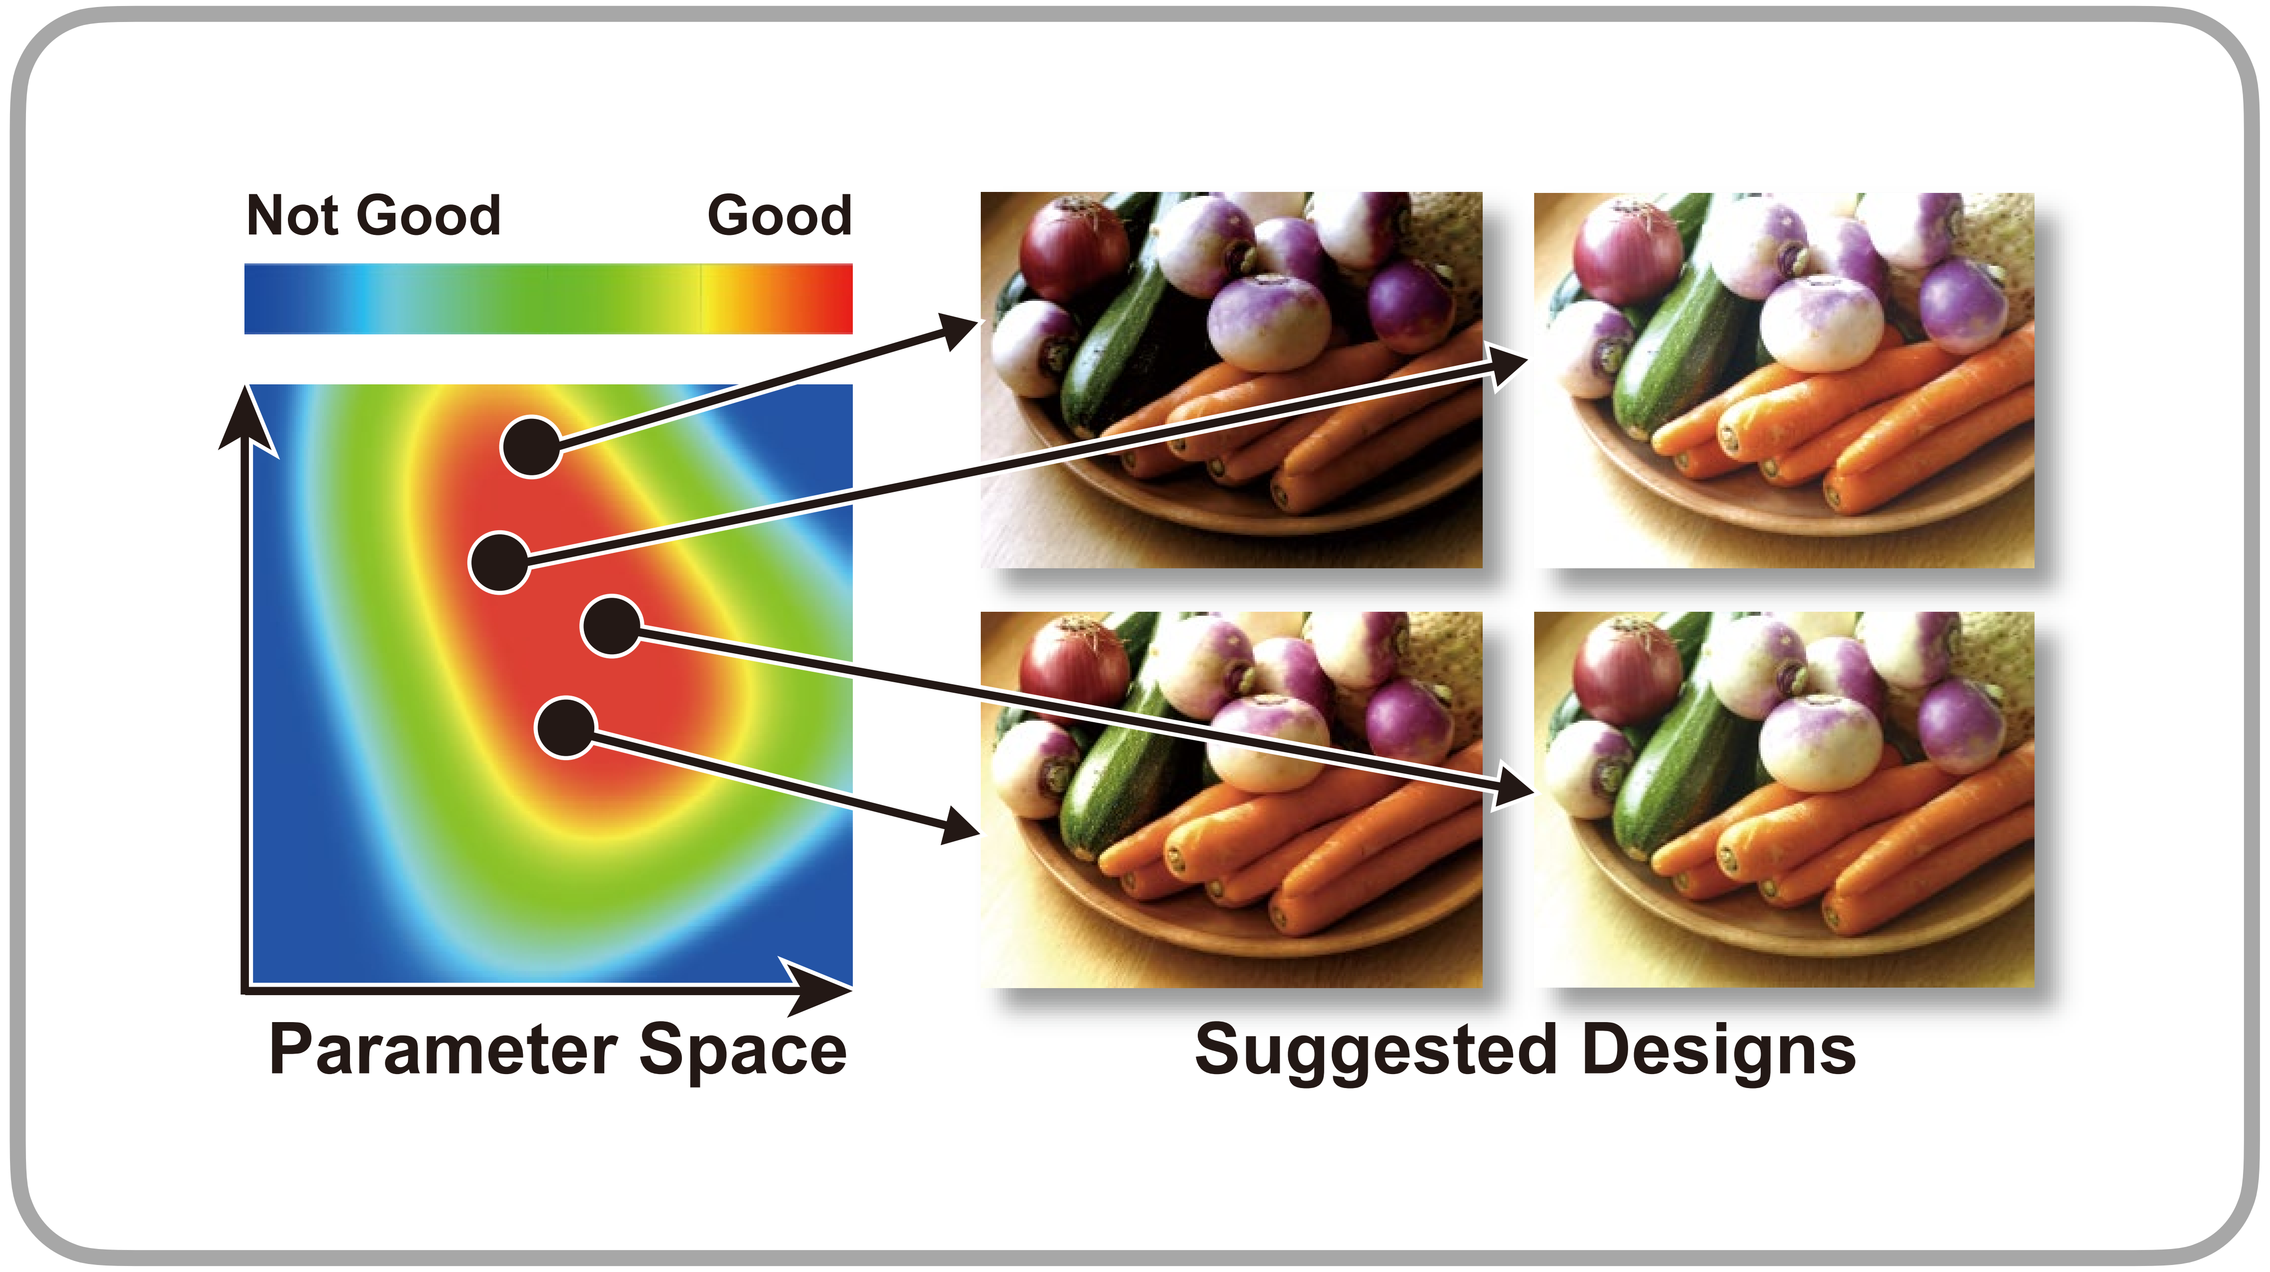

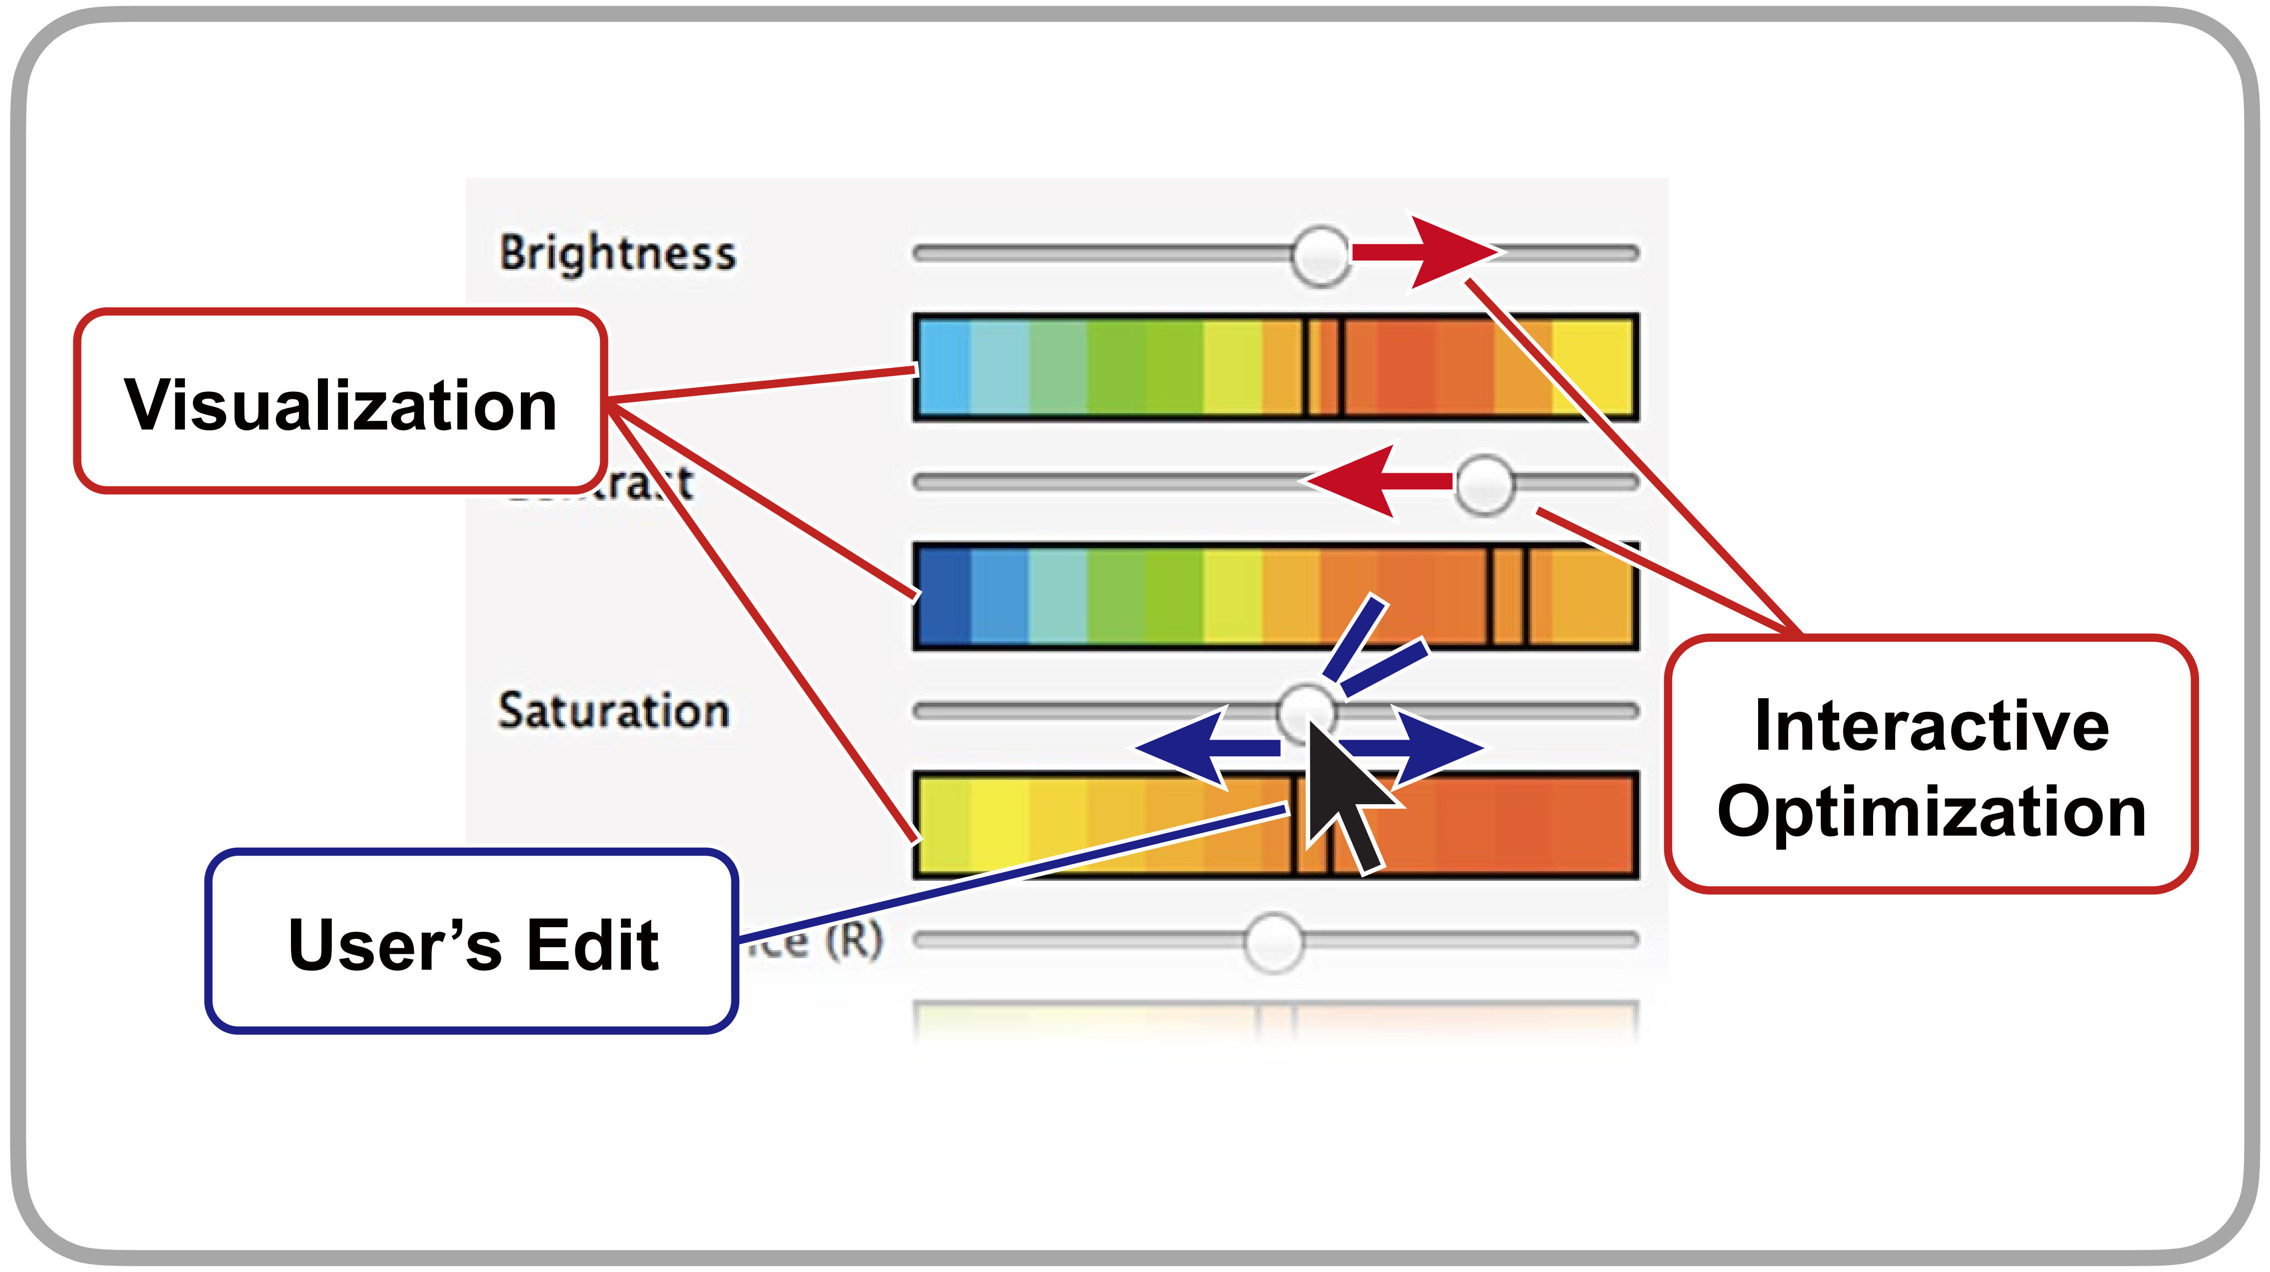

Goodness-aware **suggestions** for ideation

Goodness-aware **sliders** for guided exploration

1. Goodness-Aware Suggestions

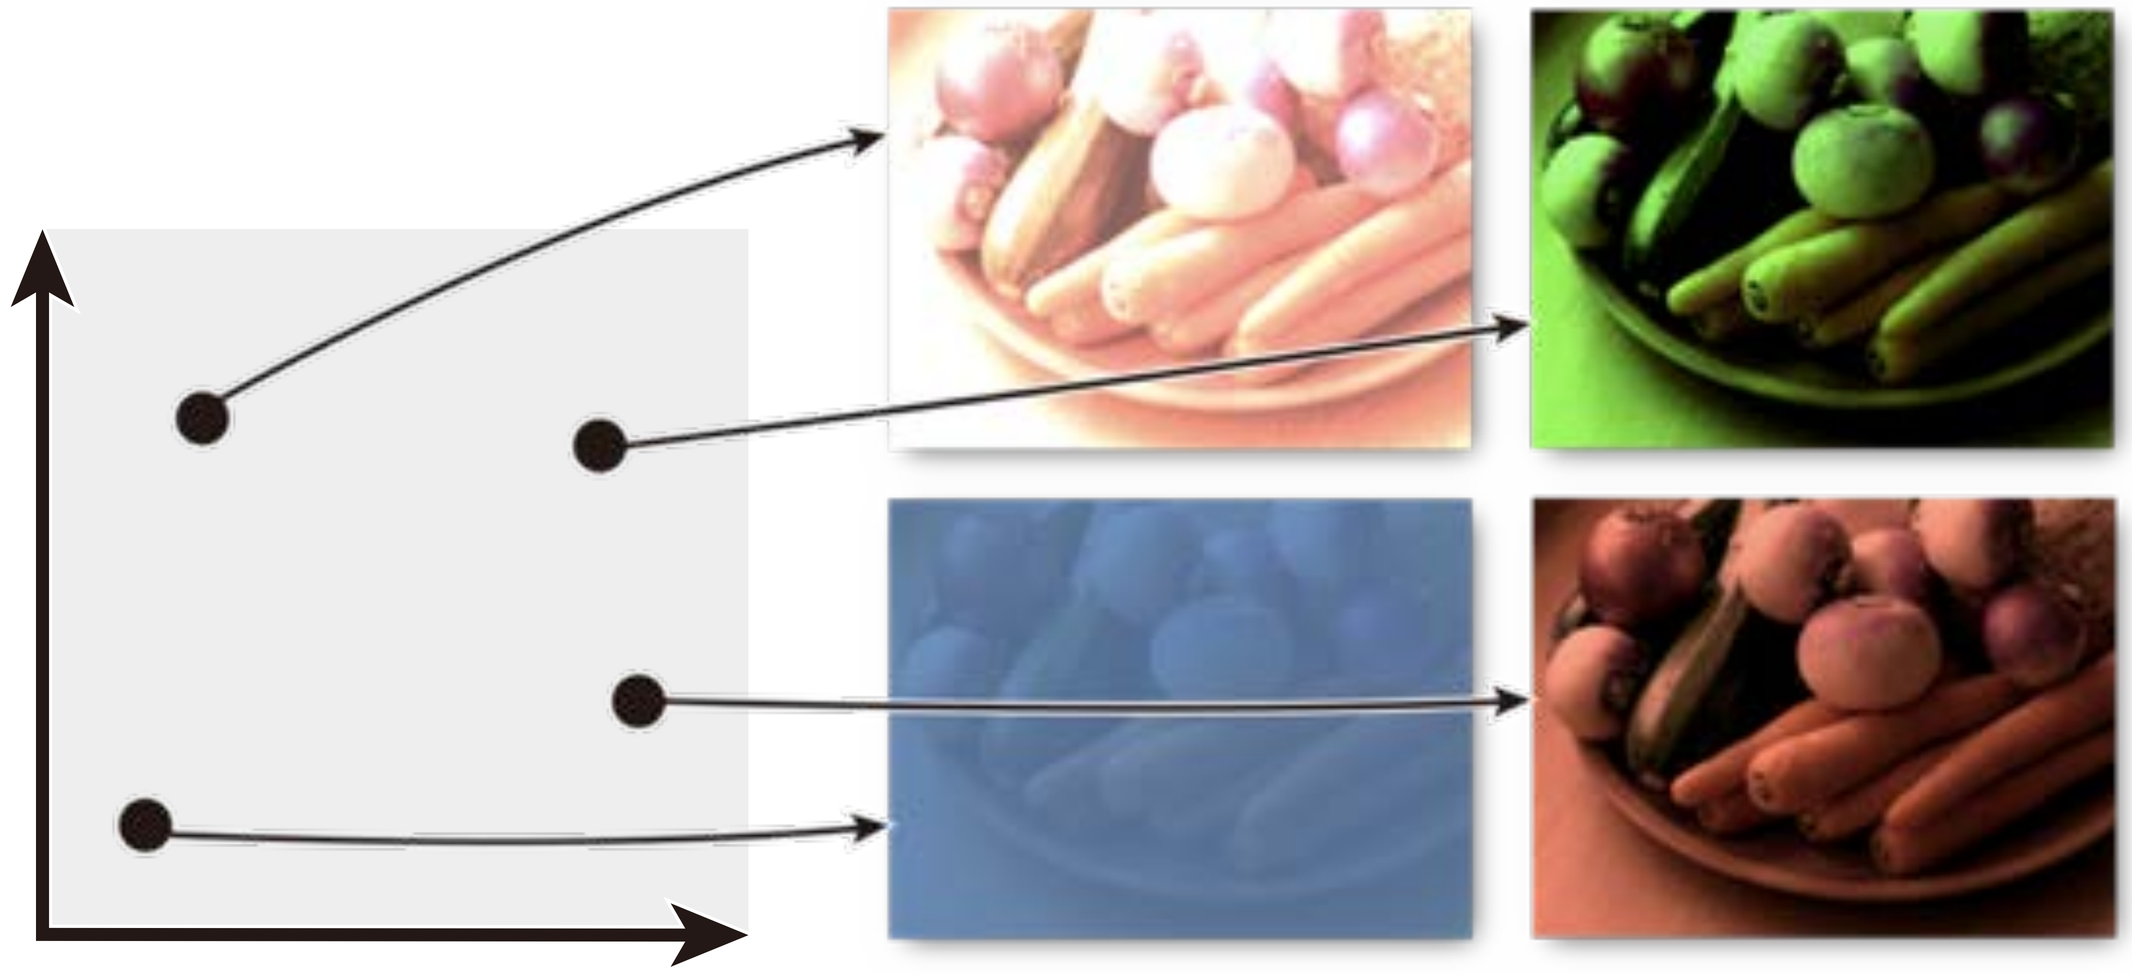

Purely random sampling (unaware of goodness) would generate many unreasonable suggestions that are not worth providing to the user

1. Goodness-Aware Suggestions

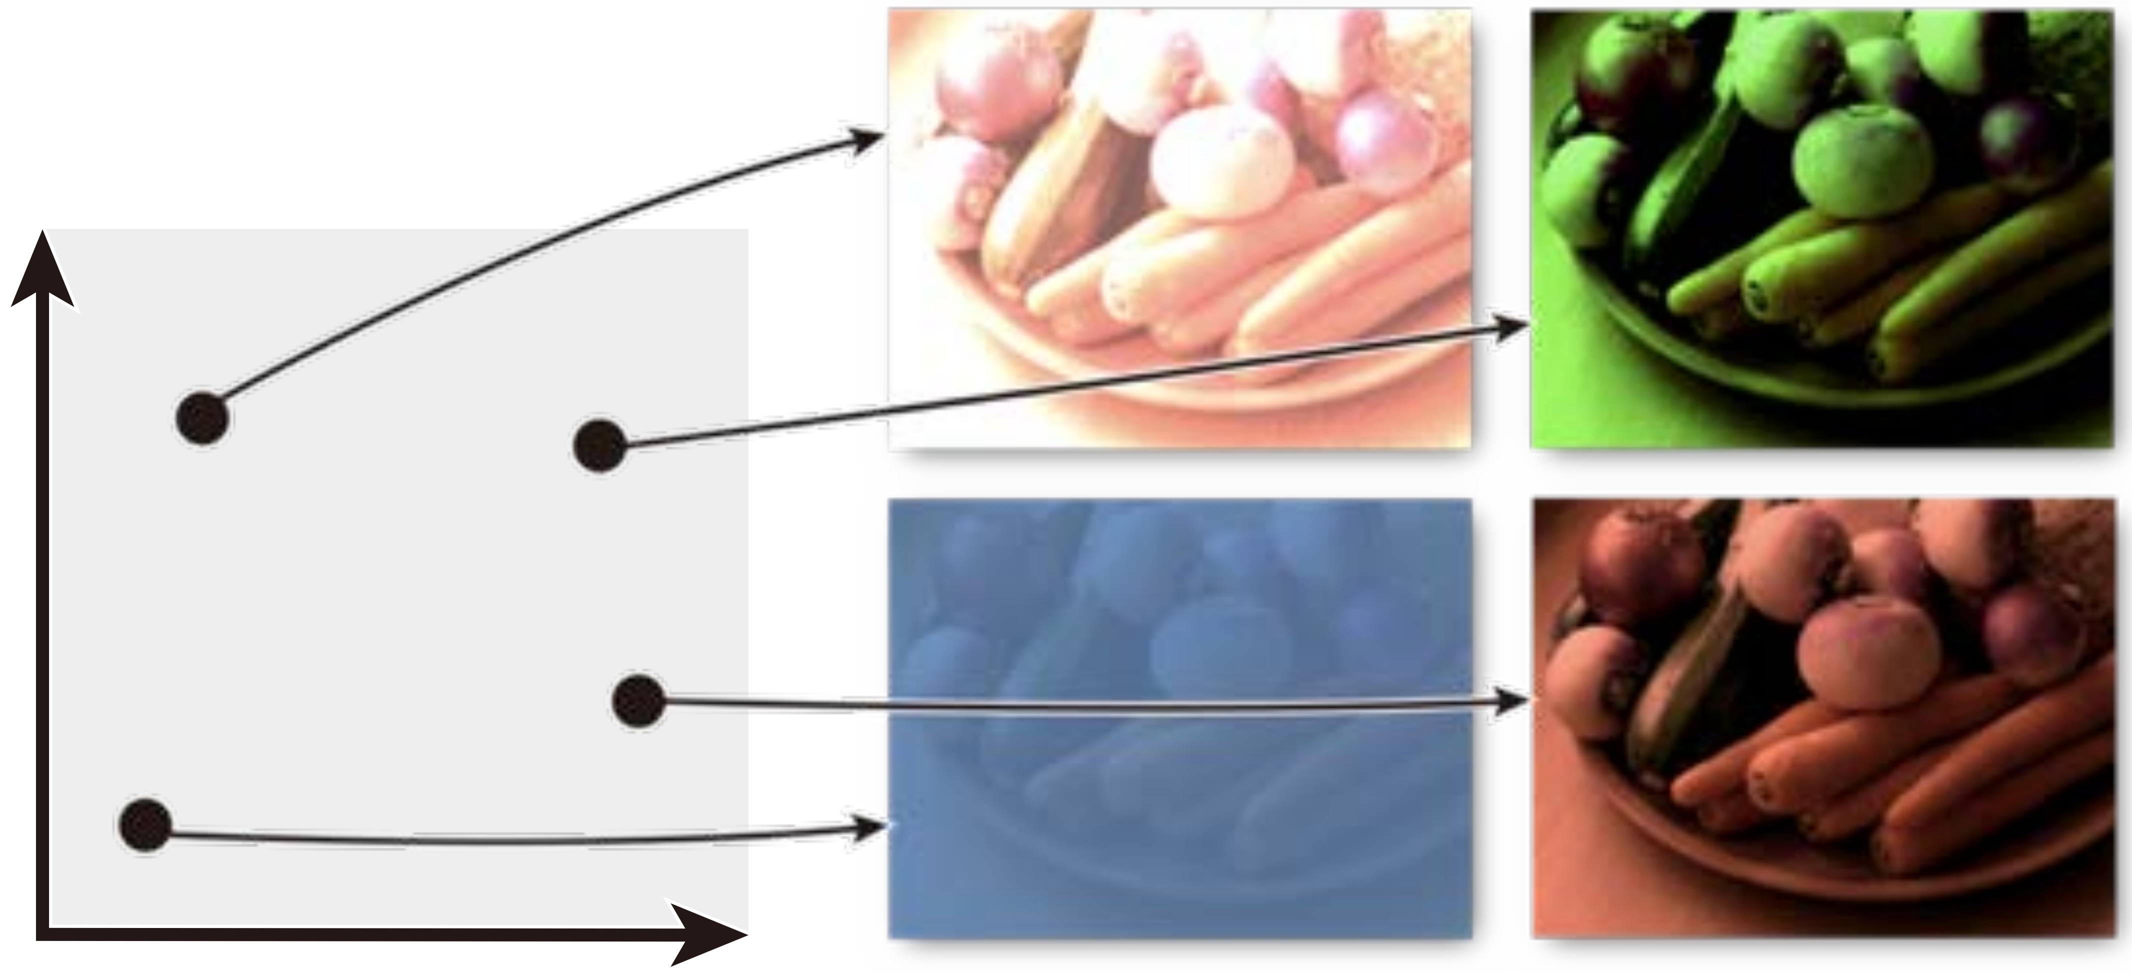

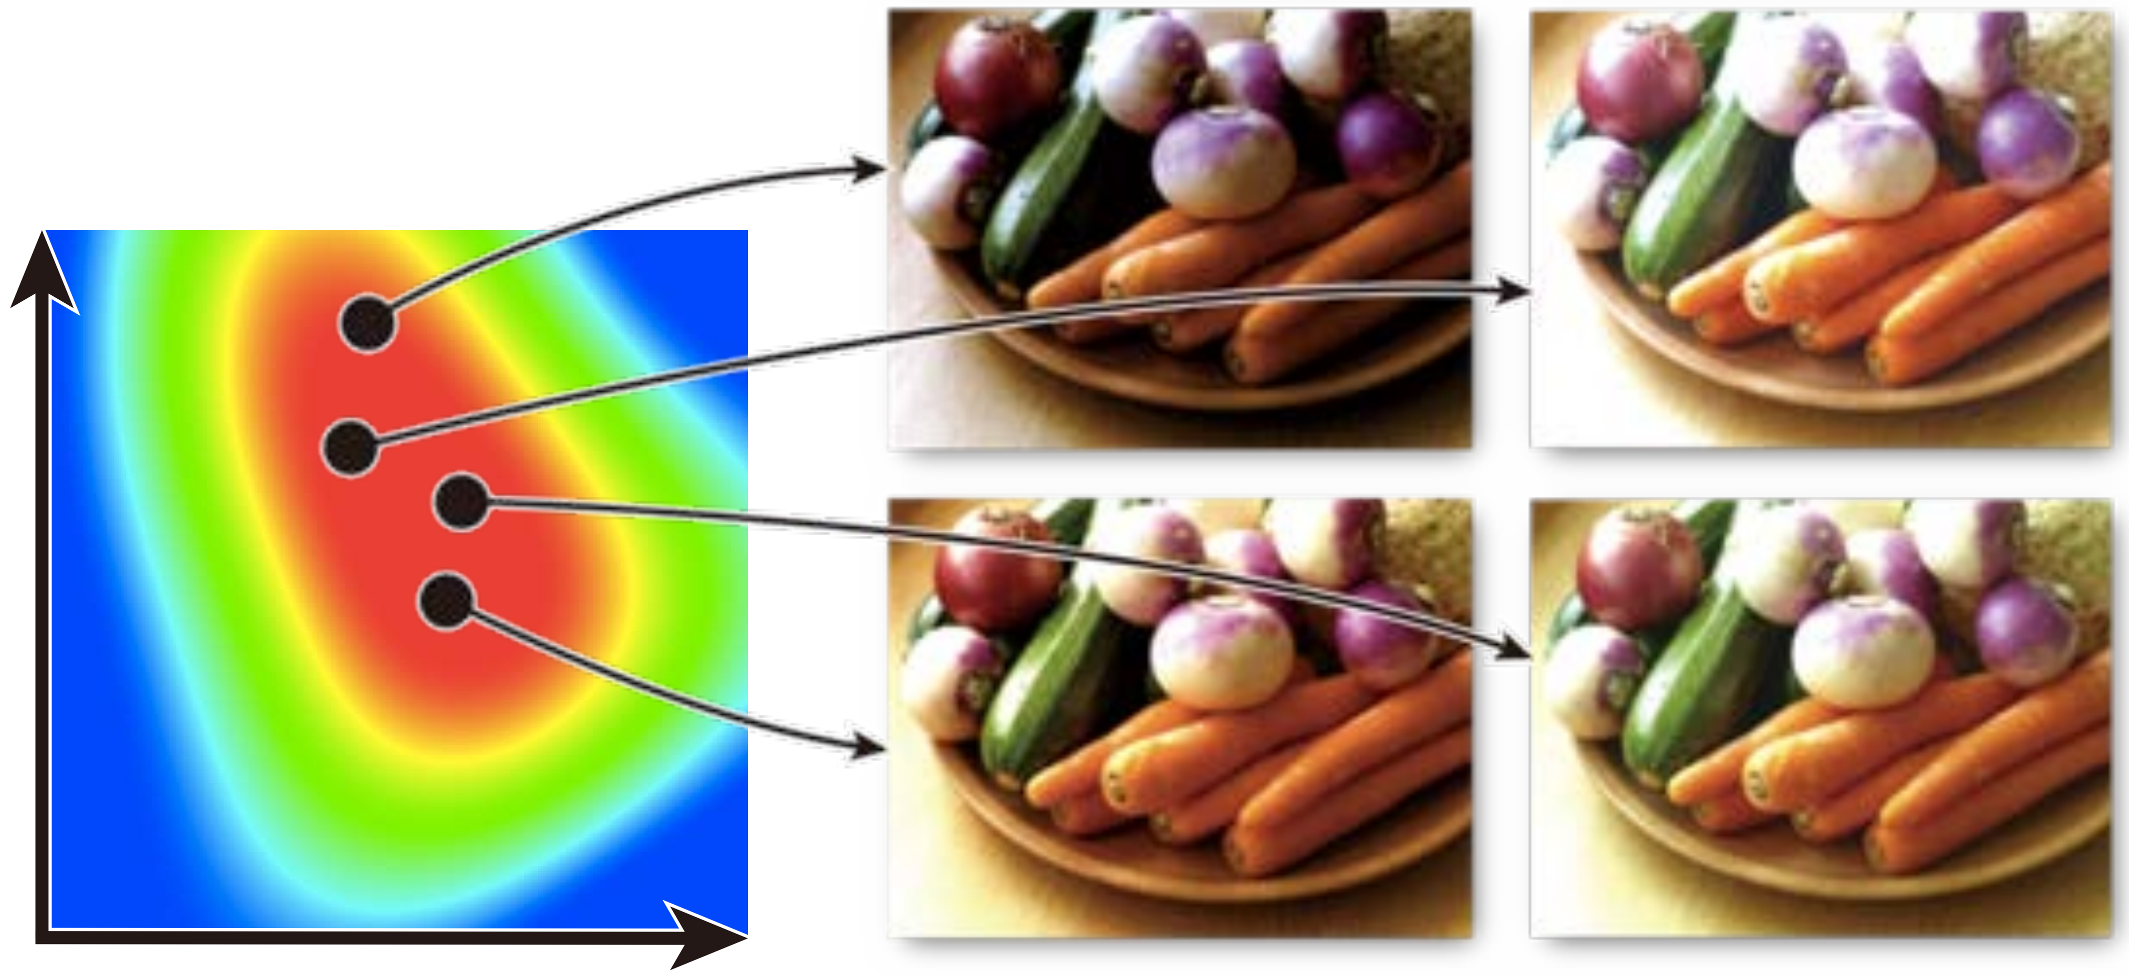

Purely random sampling (unaware of goodness) would generate many unreasonable suggestions that are not worth providing to the user

Goodness-aware sampling (biased toward the "good" region) would generate reasonable suggestions that are worth providing to the user

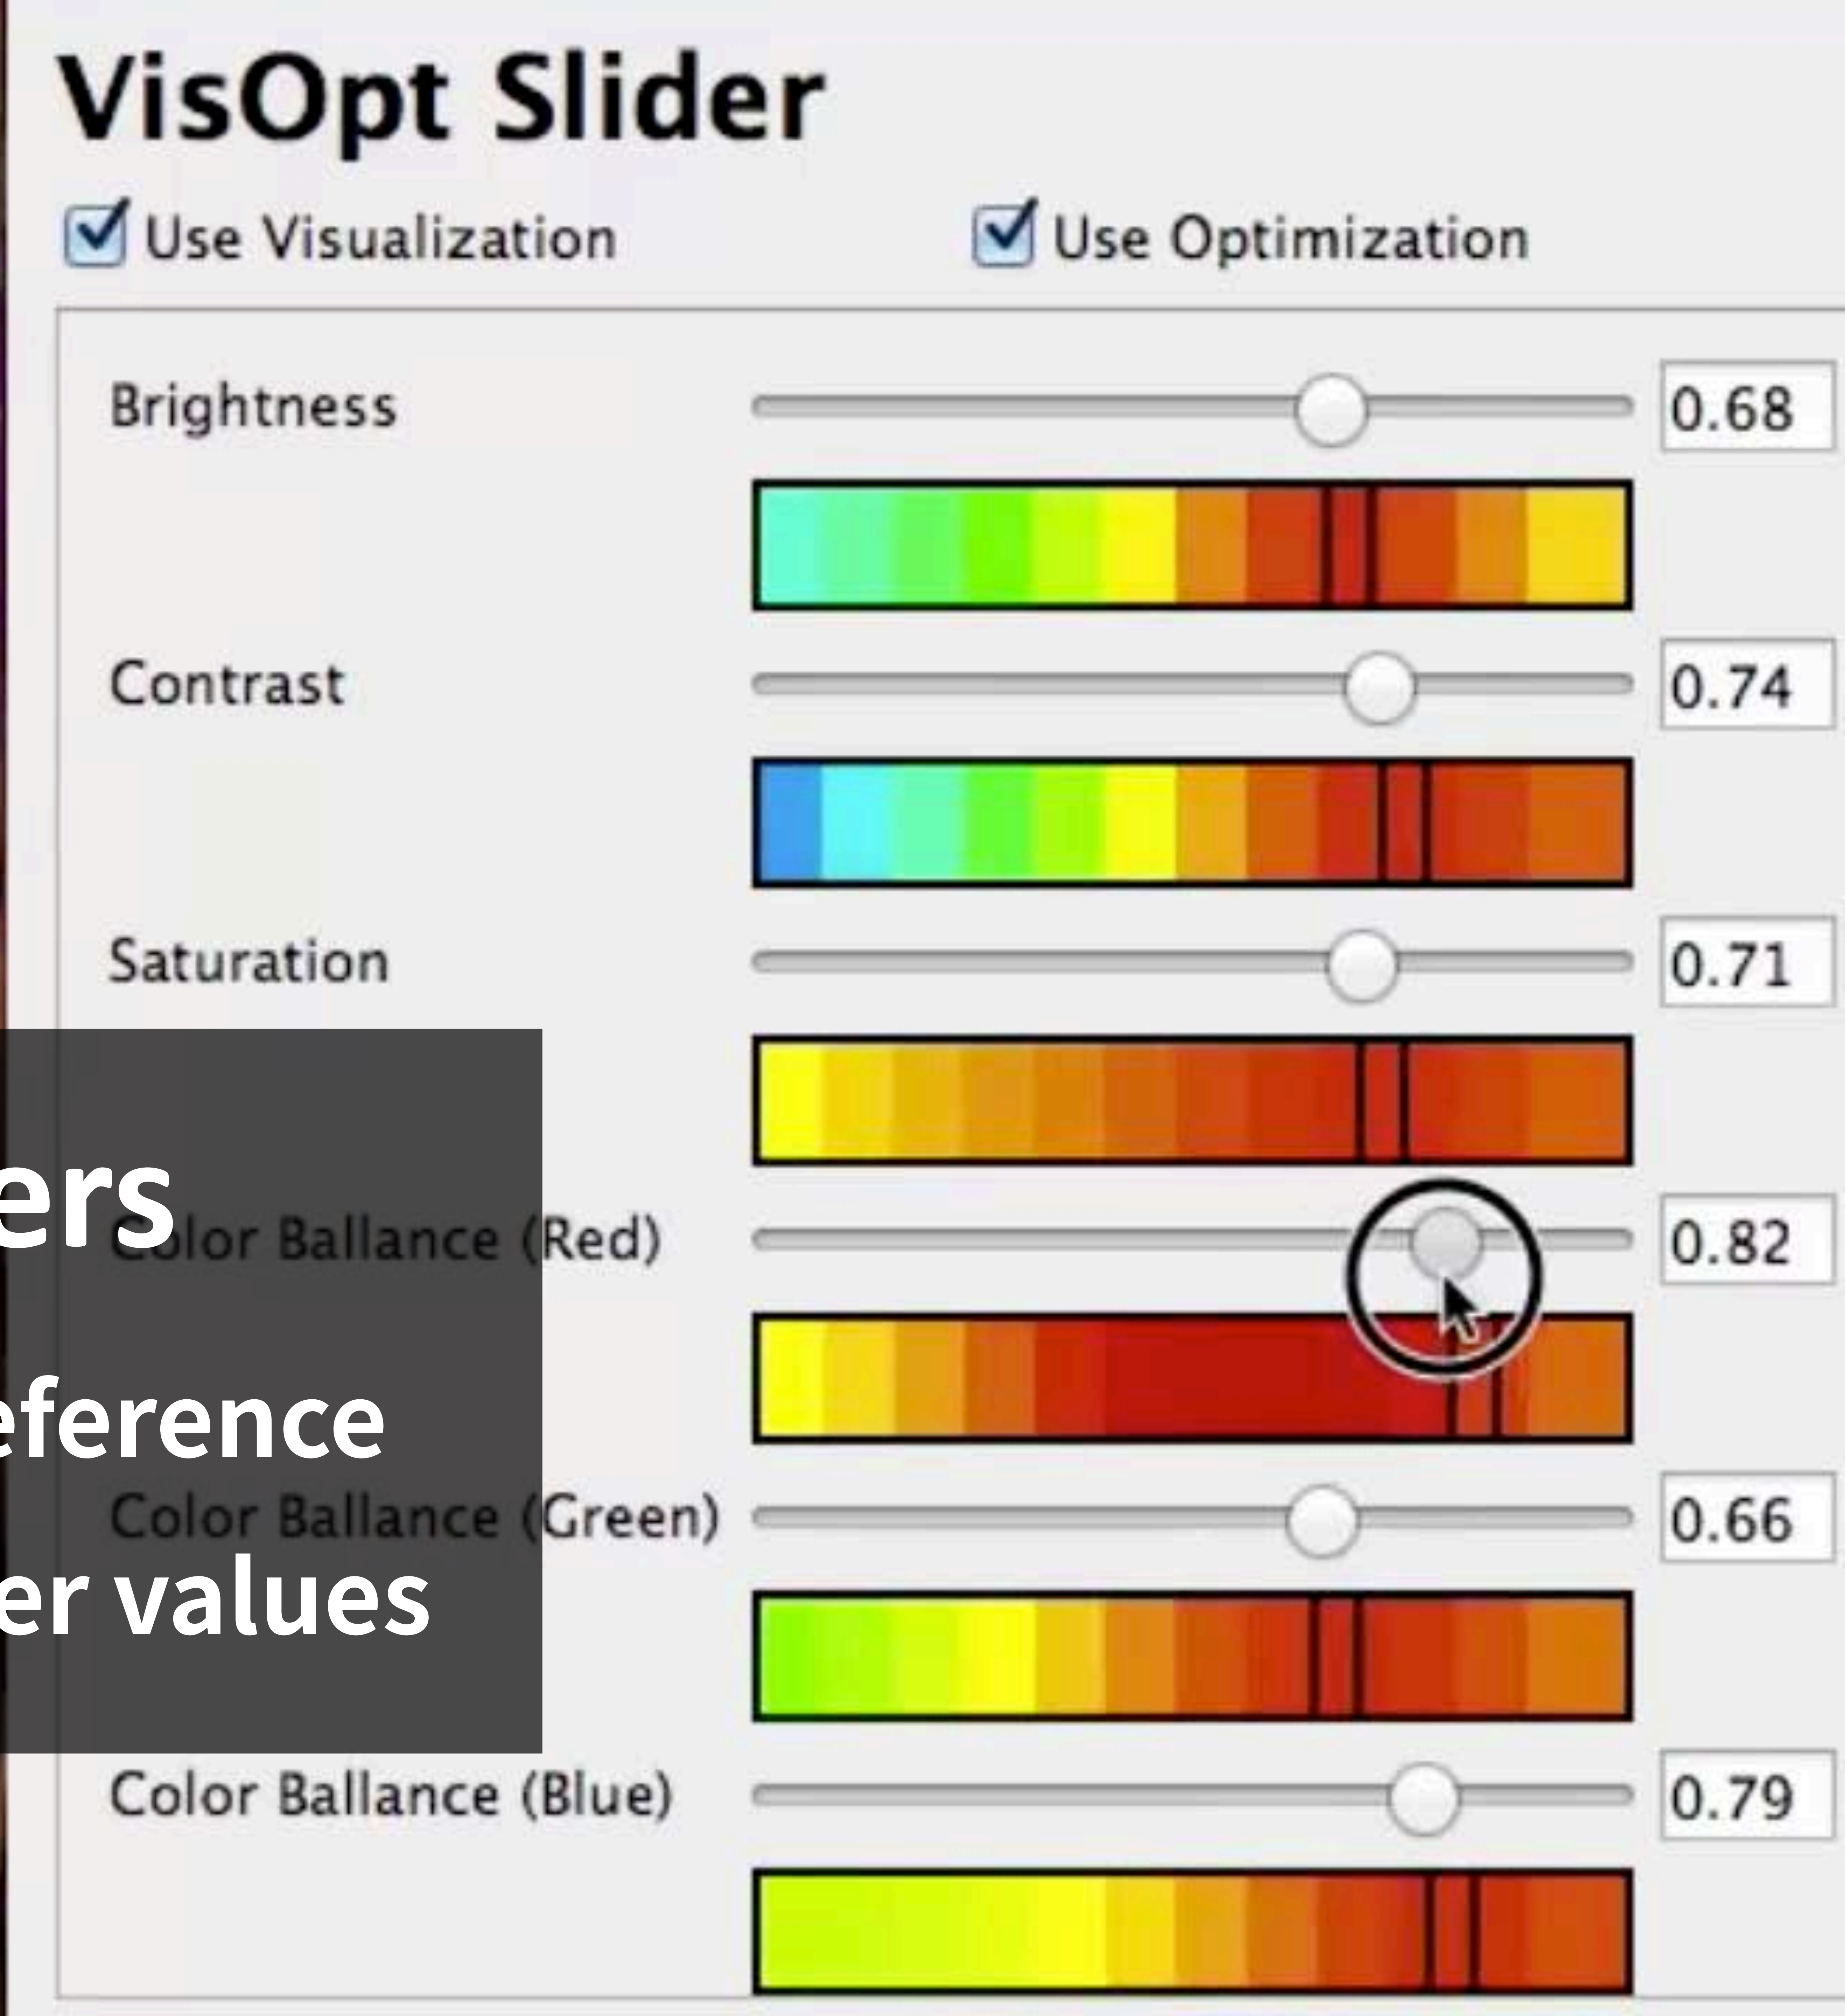

2. Goodness-Aware Sliders

- **Visualization** of the crowds' preference
- **Interactive optimization** of slider values


Design Preview
User Interface
Smart Suggestion
Make It Smart
Show Suggestions
VisOpt Slider
Use Visualization
Use Optimization
Brightness 0.74
Contrast 0.89
Saturation 0.93
Color Balance (R) 0.45
Color Balance (G) 0.58
0.72
Photo Enhancement (6D)

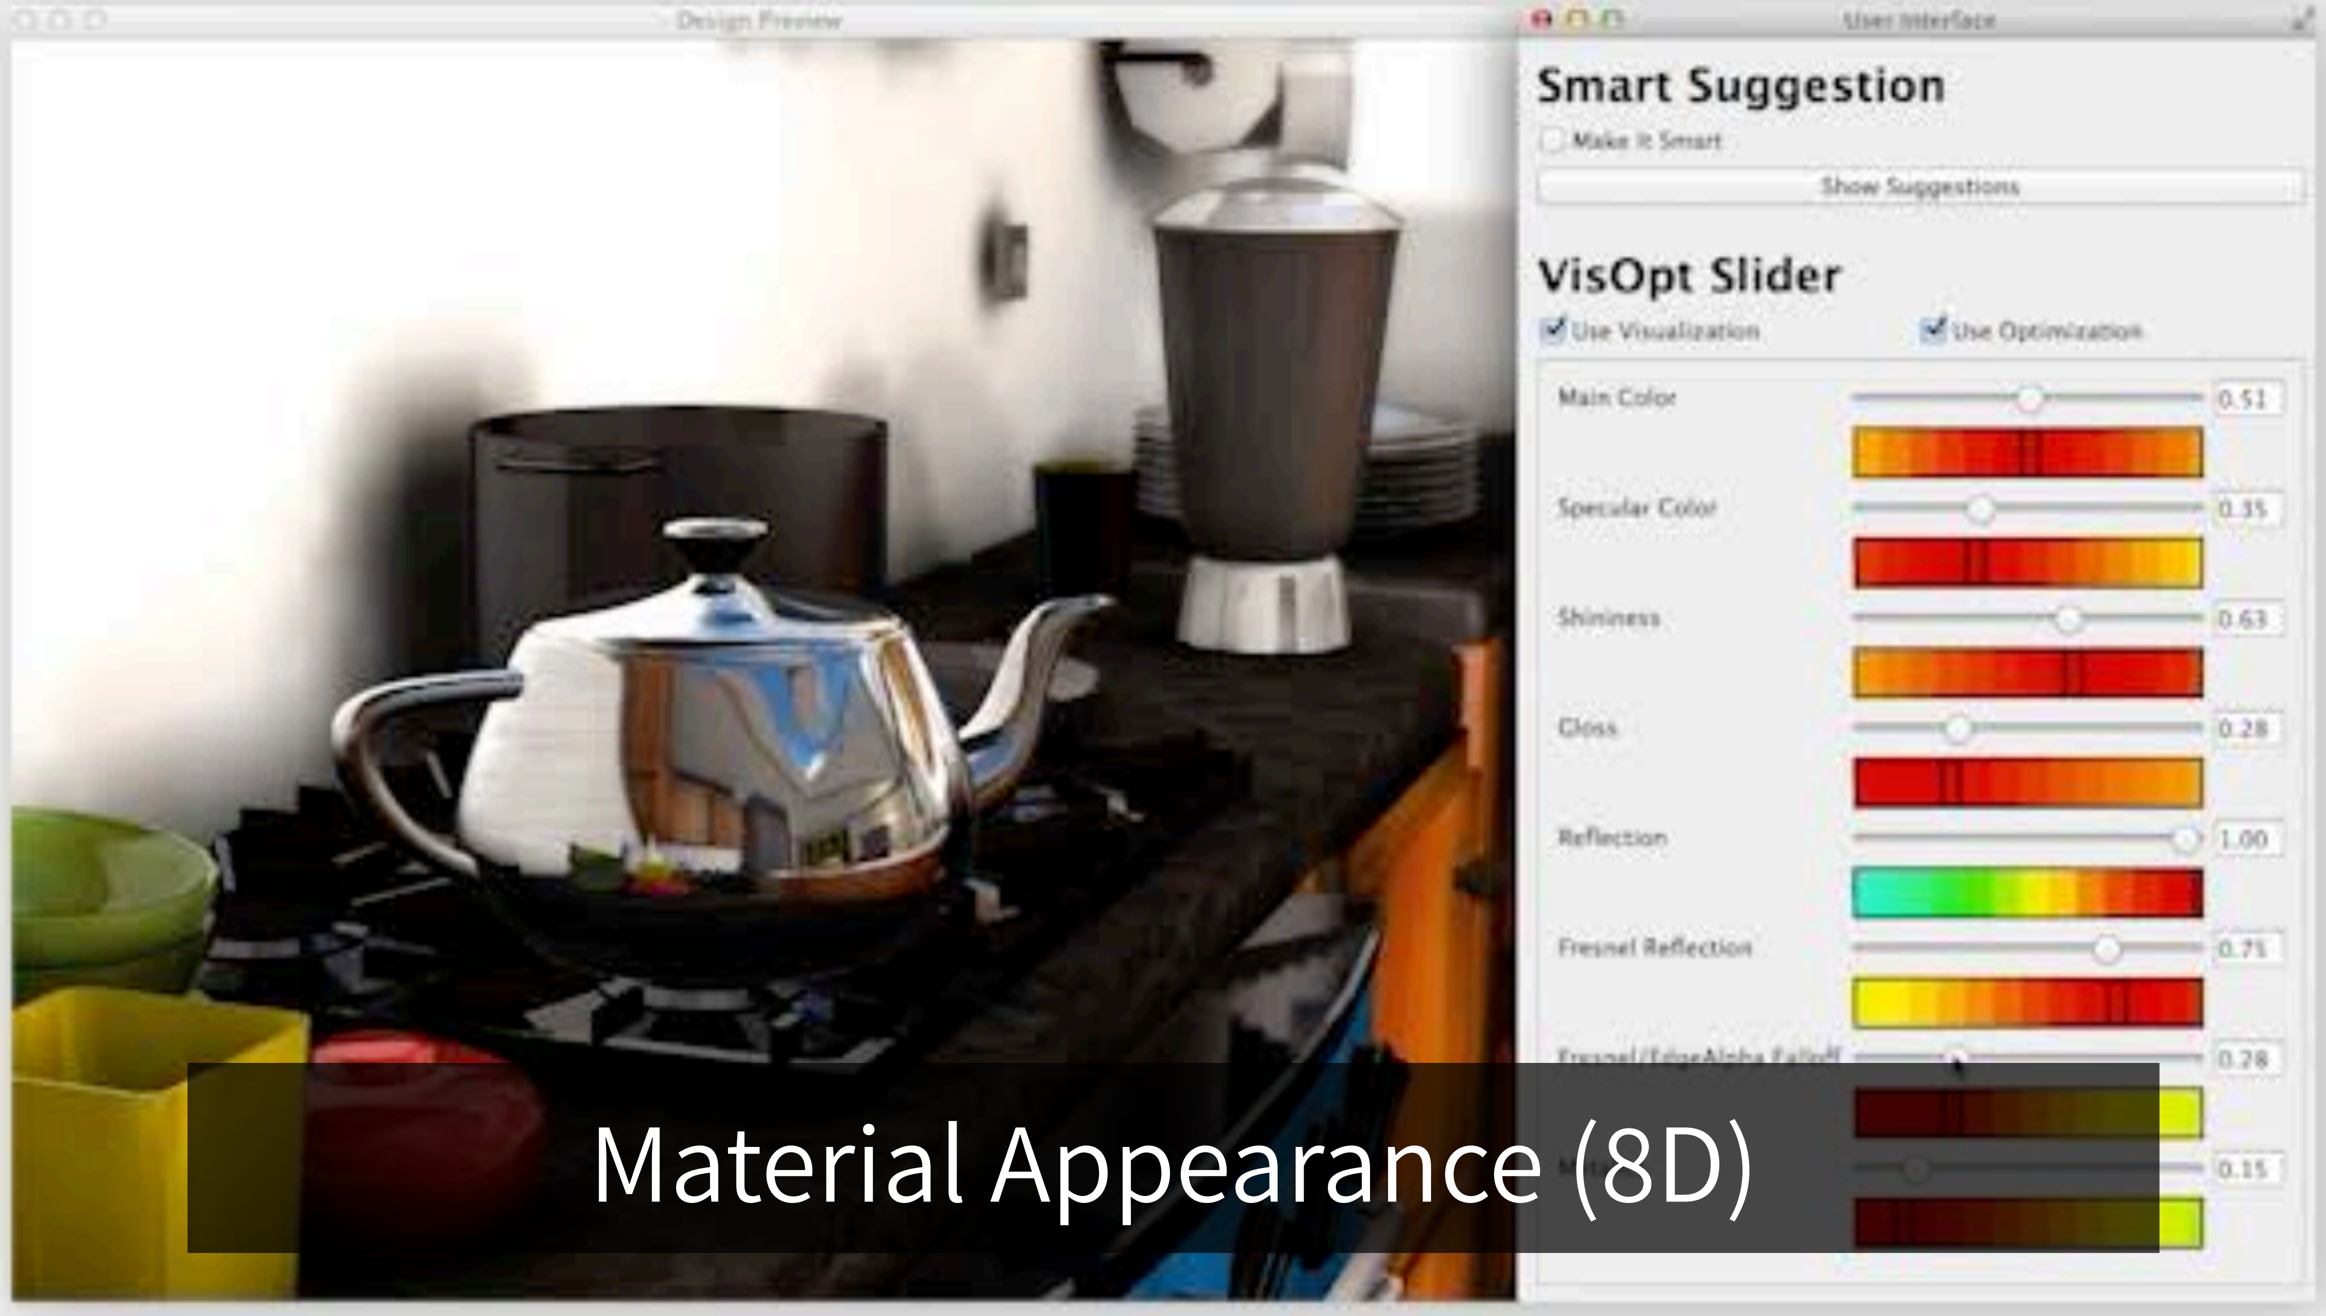
Smart Suggestion
VisOpt Slider
Material Appearance (8D)

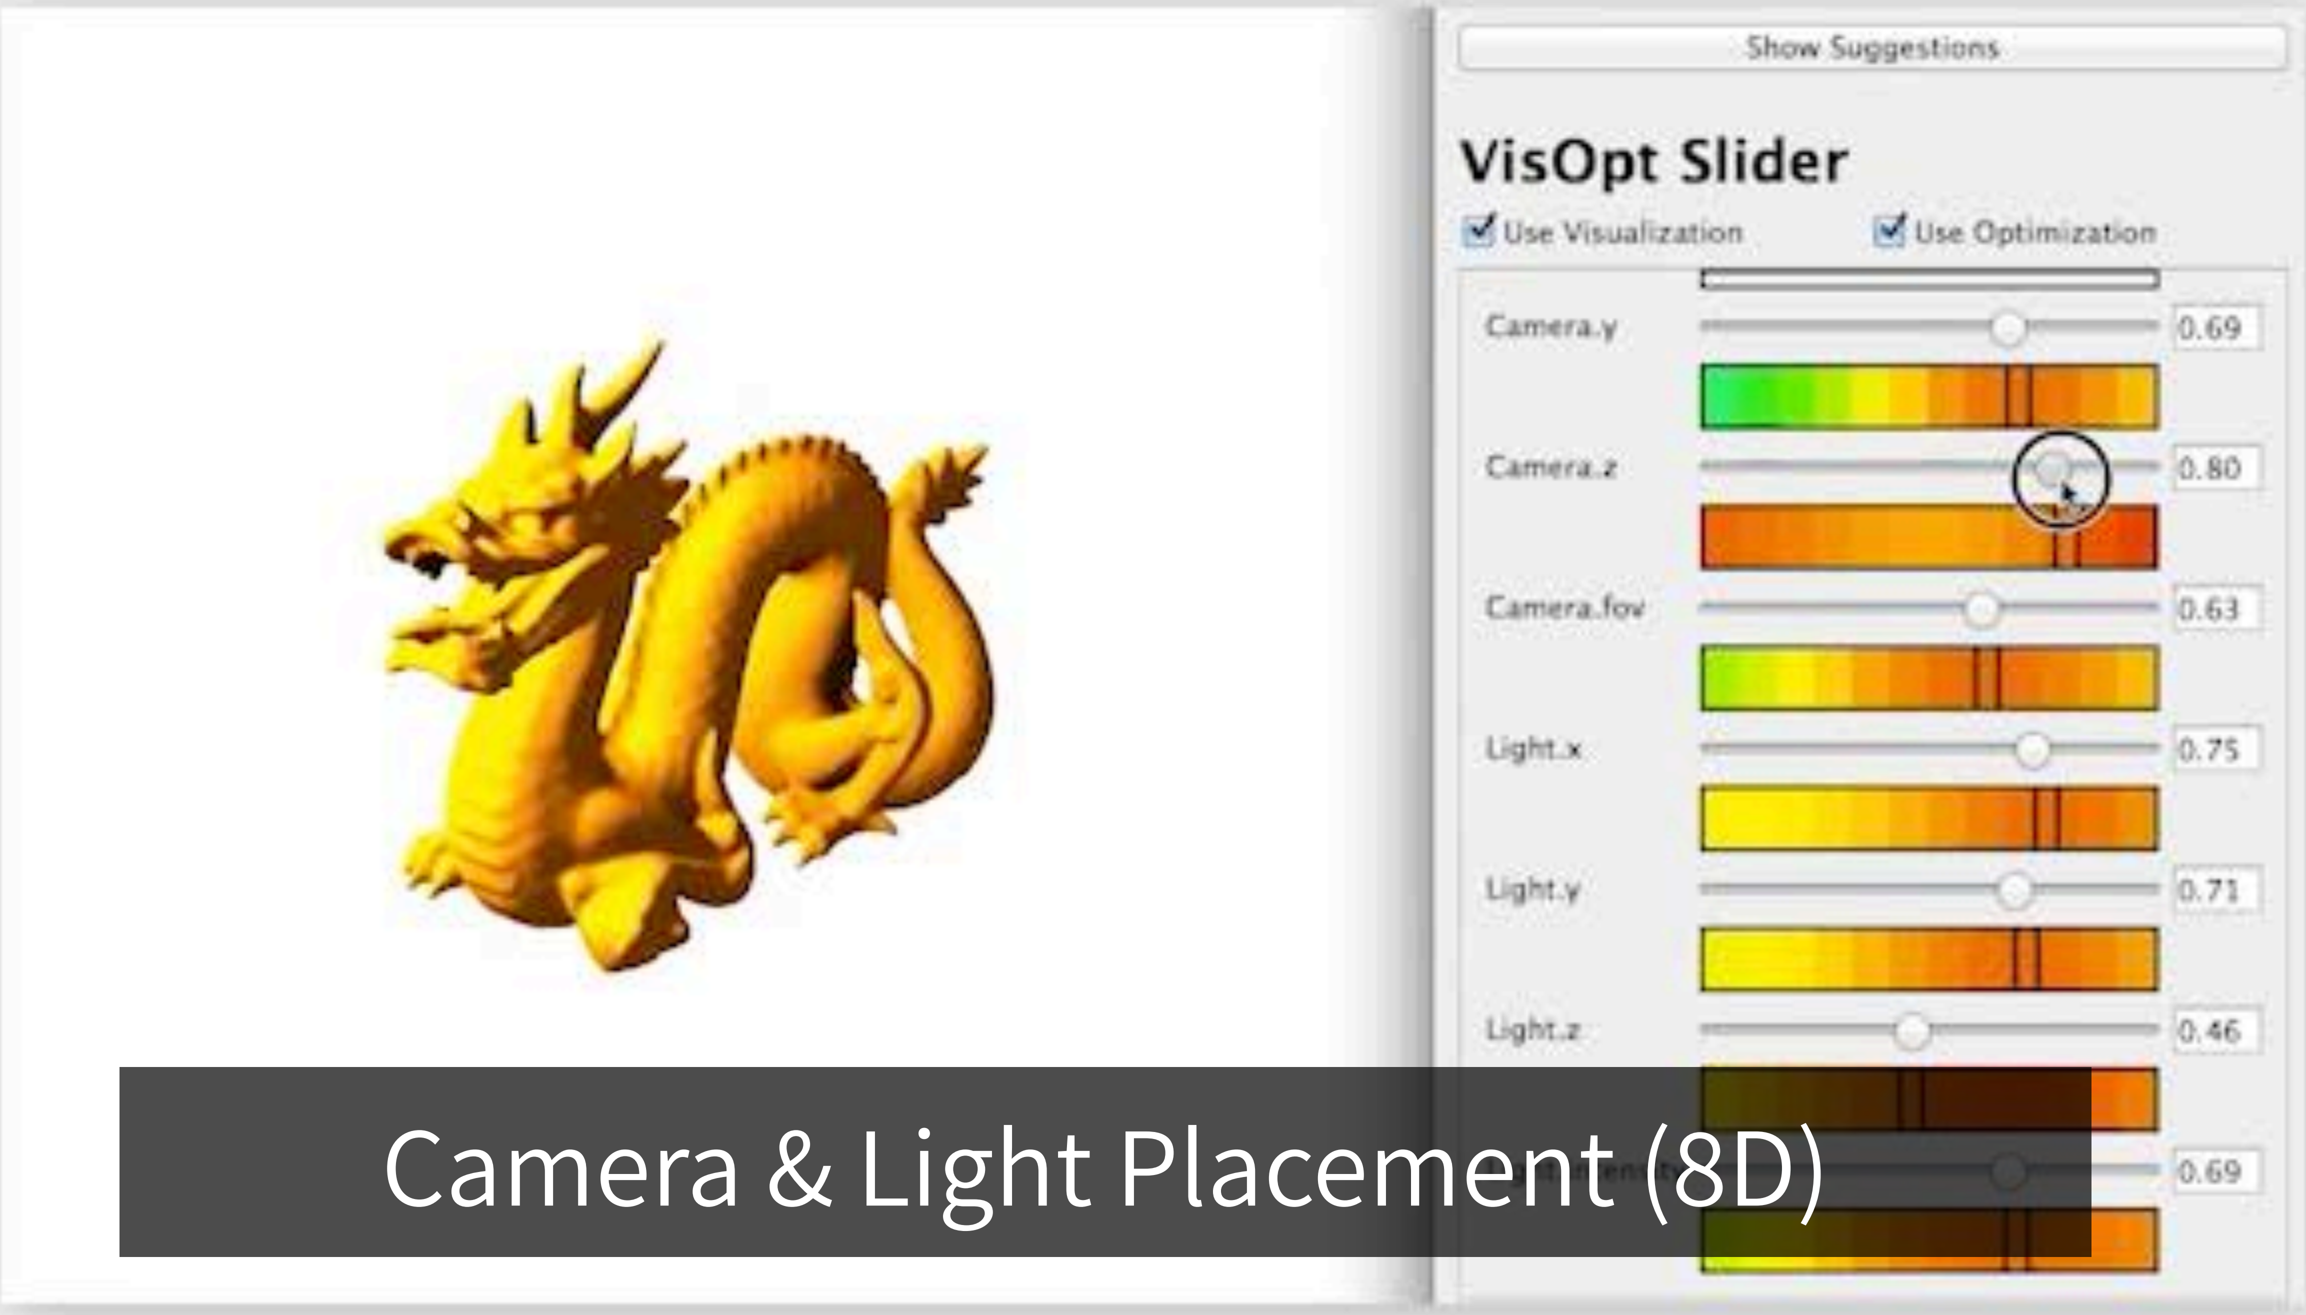
Show Suggestions
VisOpt Slider
Use Visualization
Use Optimization
Camera.y 0.69
Camera.z 0.80
Camera.fov 0.63
Light.x 0.75
Light.y 0.71
Light.z 0.46
0.69
Camera & Light Placement (8D)


Smart Suggestion
VisOpt Slider
Facial Expression (53D)

Intelligent Tools Case 2

Intelligent Automatic Solver

[Koyama+, SIGGRAPH 2017] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. ACM Trans. Graph. 36, 4, pp.48:1–48:11 (2017). https://doi.org/10.1145/3072959.3073598

“Crowdsource” button in the tool

“People’s Choice” optimal slider values

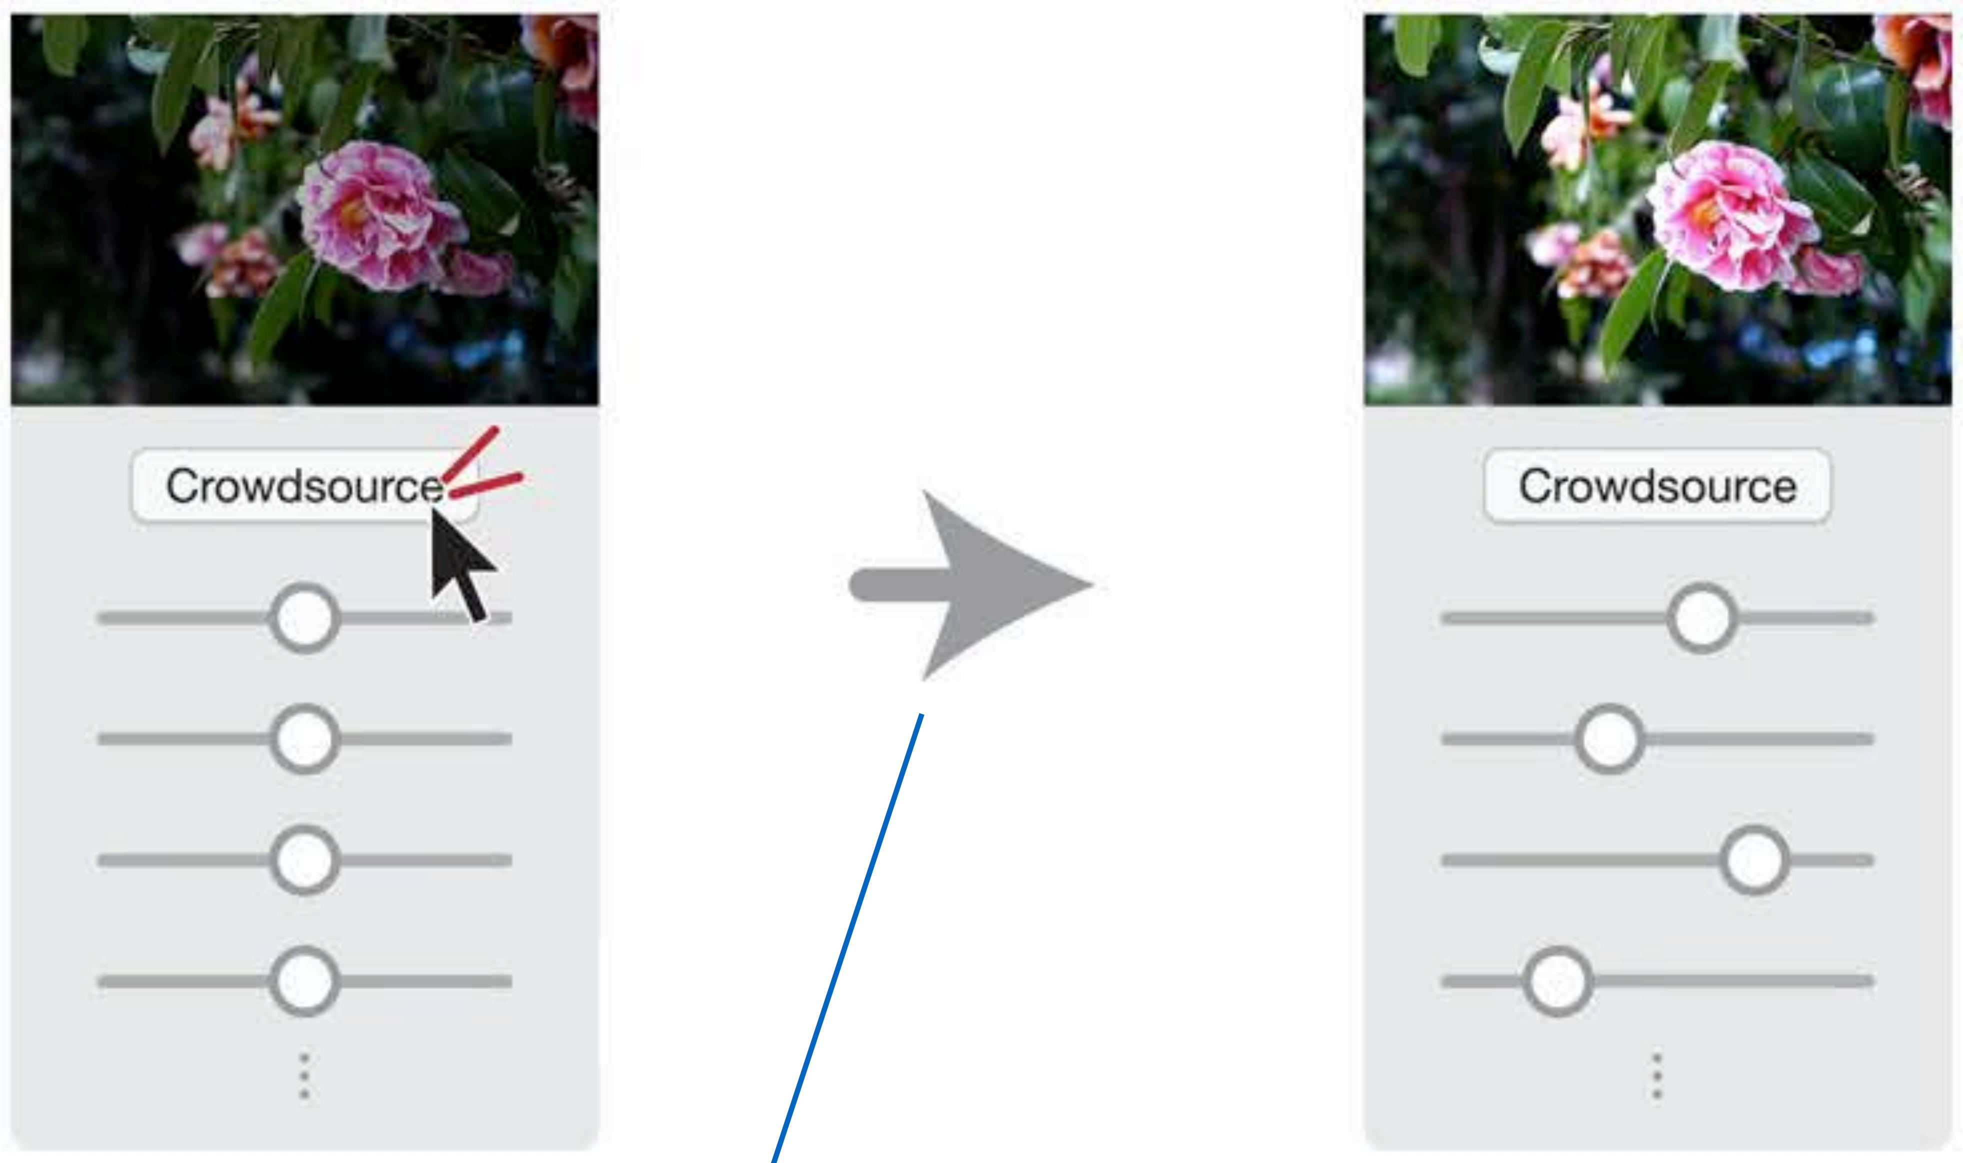

Crowd-powered optimization

[Koyama+, SIGGRAPH 2017] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. ACM Trans. Graph. 36, 4, pp.48:1–48:11 (2017). https://doi.org/10.1145/3072959.3073598

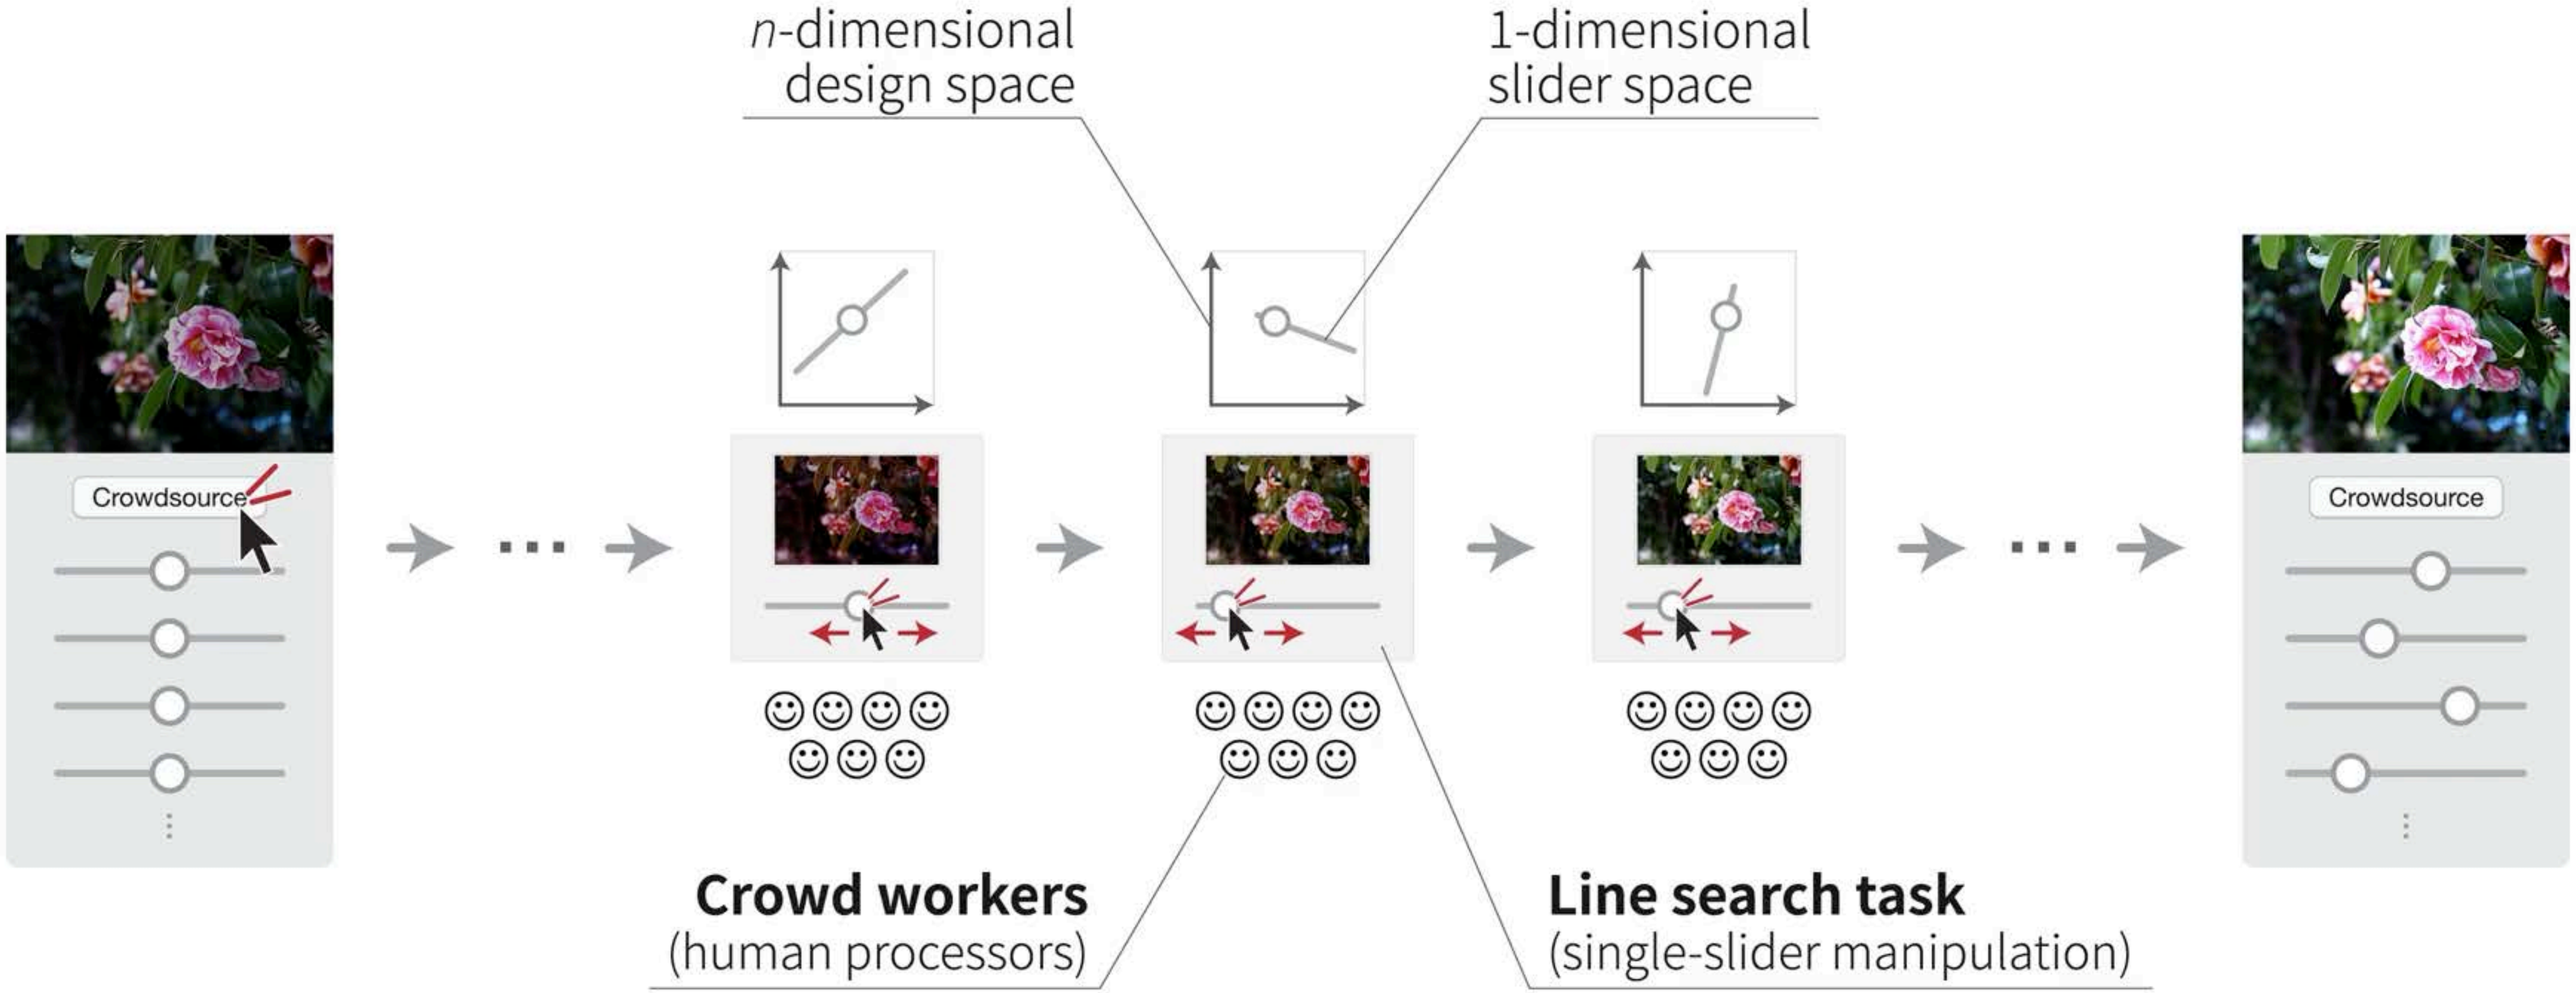

[Koyama+, SIGGRAPH 2017] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. ACM Trans. Graph. 36, 4, pp.48:1–48:11 (2017). https://doi.org/10.1145/3072959.3073598

“Relative Assessment” Microtask Design

Task: Choose the image that looks better

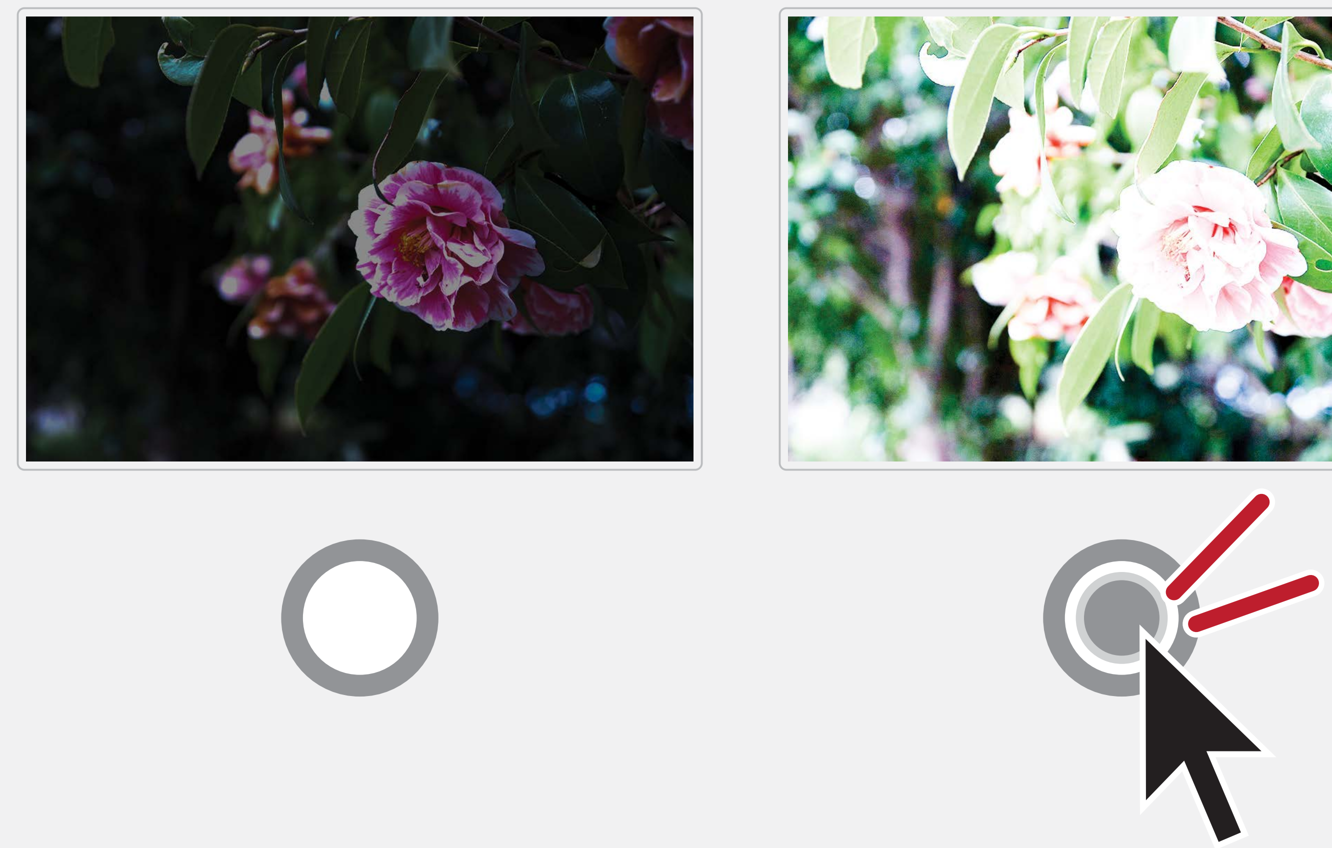

Task: Adjust the slider so that the image looks the best

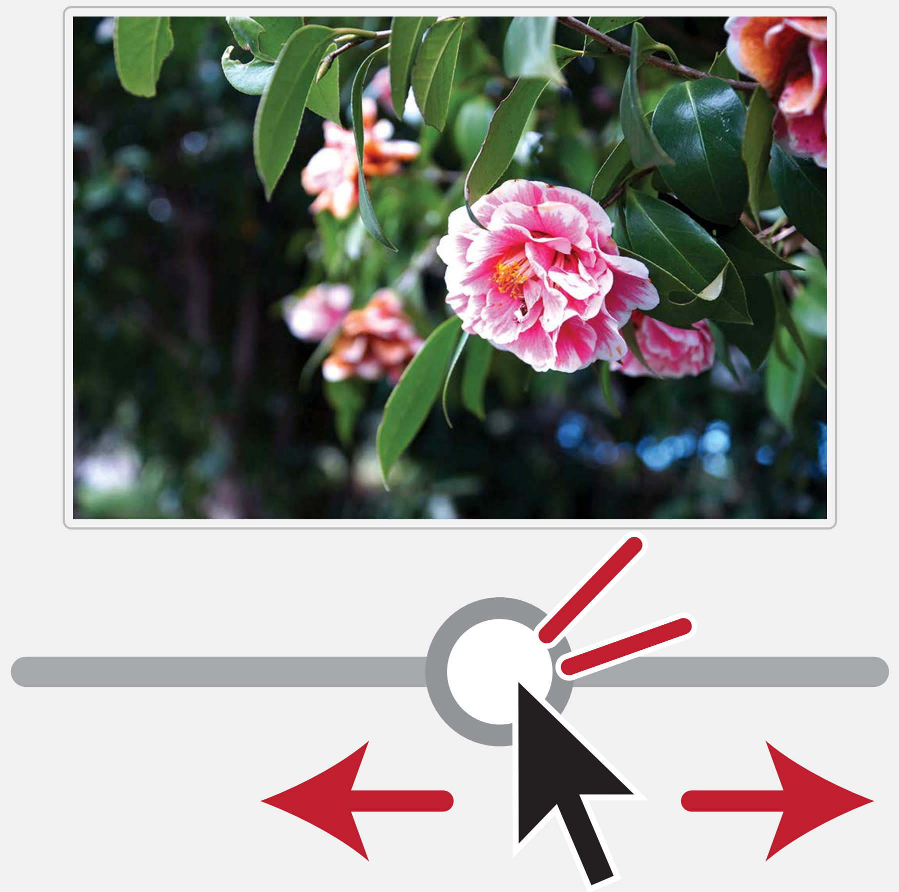

Basic: Pairwise comparison
(e.g., [Koyama+, UIST 2014])

Advanced: Single-slider manipulation
(provides much richer information)

➡ Faster convergence

Adjust the slider so that the image looks best.

Web interface for crowdsourcing

No.2

Adjust the slider so that the image looks best.

Algorithm: Sequential Line Search

A human-in-the-loop Bayesian optimization

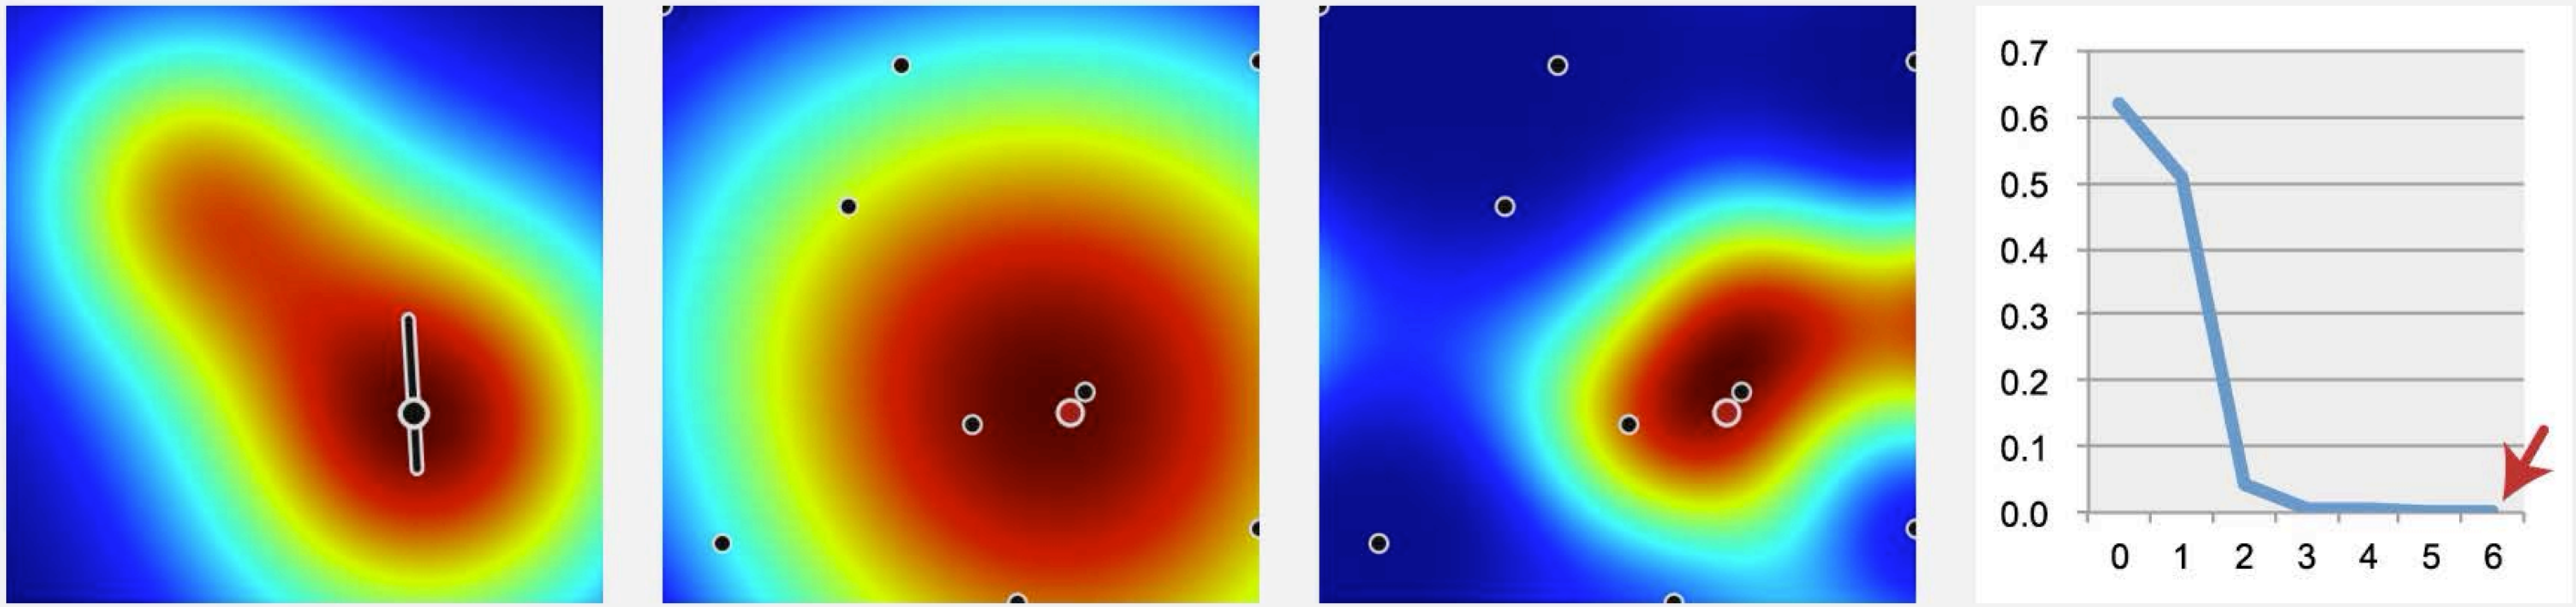

Objective function | Estimated function | Expected improvement | Residuals over iterations

Please refer to [Koyama+, SIGGRAPH 2017] for the mathematical details

[Koyama+, SIGGRAPH 2017] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. ACM Trans. Graph. 36, 4, pp.48:1–48:11 (2017). https://doi.org/10.1145/3072959.3073598

Applications #1

Photo Color Enhancement (6D)

Original Photographs

For each photo, it runs **15 iterations**, cost **5.25 USD** in total, and took **68 min** in average

Results






For each photo, it runs **15 iterations**, cost **5.25 USD** in total, and took **68 min** in average

Replay

Original Photographs

For each photo, it runs **15 iterations**, cost **5.25 USD** in total, and took **68 min** in average

Results






For each photo, it runs **15 iterations**, cost **5.25 USD** in total, and took **68 min** in average

Evaluation: Crowdsourced Voting

Q. Which one do you like?

Original	By Crowds	By Photoshop	By Lightroom
			

	Original	Crowds	Photoshop	Lightroom
				
Preferred by:	2	26	2	3
				
Preferred by:	0	32	0	1

	Original	Crowds	Photoshop	Lightroom
				
Preferred by:	0	29	1	3
				
Preferred by:	1	23	3	6

	Original	Crowds	Photoshop	Lightroom
		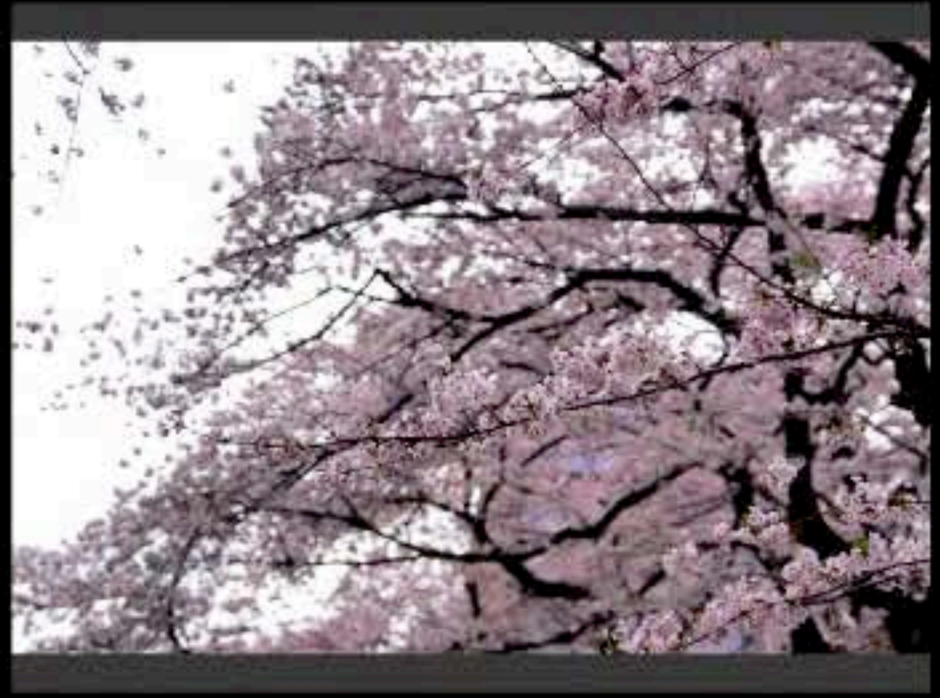		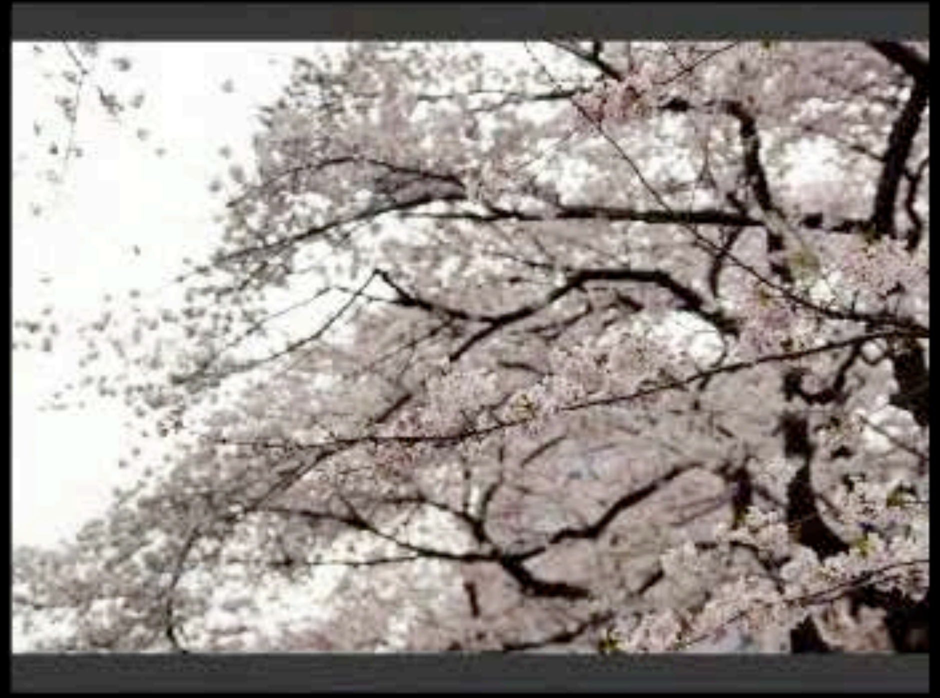
Preferred by:	0	29	3	1
				
Preferred by:	0	31	0	2

Q. Which one do you like?

Original

By Crowds

By Photoshop

By Lightroom

➡ As they are the “people’s choice” enhancements, people preferred them

Applications #2

Material Appearance (3D / 7D)

Usage Scenario #1

Material Appearance from Reference Photograph

Reference

Result (3D)

Reference

Result (3D)

Usage Scenario #2

Material Appearance from Text Instruction

Input instruction:

“Mirror-like reflective”

Result (3D)

Input instruction:

“Brushed strainless”

Result (3D)

Discussions

Other Types of Intelligence

Types of Intelligence Enabled by Crowdsourcing

- **General preference**
  - [Koyama+, UIST 2014]
  - [Koyama+, SIGGRAPH 2017]
  - [Secord+, TOG (2012)]
  - [Zhu+, SA 2014]

Please note that this list is not exhaustive

Types of Intelligence Enabled by Crowdsourcing

- **General preference**
  - [Koyama+, UIST 2014]
  - [Koyama+, SIGGRAPH 2017]
  - [Secord+, TOG (2012)]
  - [Zhu+, SA 2014]
- **Semantic attributes**
  - [Yumer+, SIGGRAPH 2015]
  - [O'Donovan+, SIGGRAPH 2014]
  - [Chaudhuri+, UIST 2013]
- **Perceptual similarity**
  - [Garces+, SIGGRAPH 2014],
  - [O'Donovan+, SIGGRAPH 2014]
- **Perceptual compatibility**
  - [Laursen+, PG Short 2016]
  - [Liu+, SIGGRAPH 2015]
- **… etc.**

Please note that this list is not exhaustive

Types of Intelligence [1/4]: General Preference

E.g., by learning crowds' preference in portrait expressions, the tool can intelligently select best expressions from a long video or a large collection of photographs

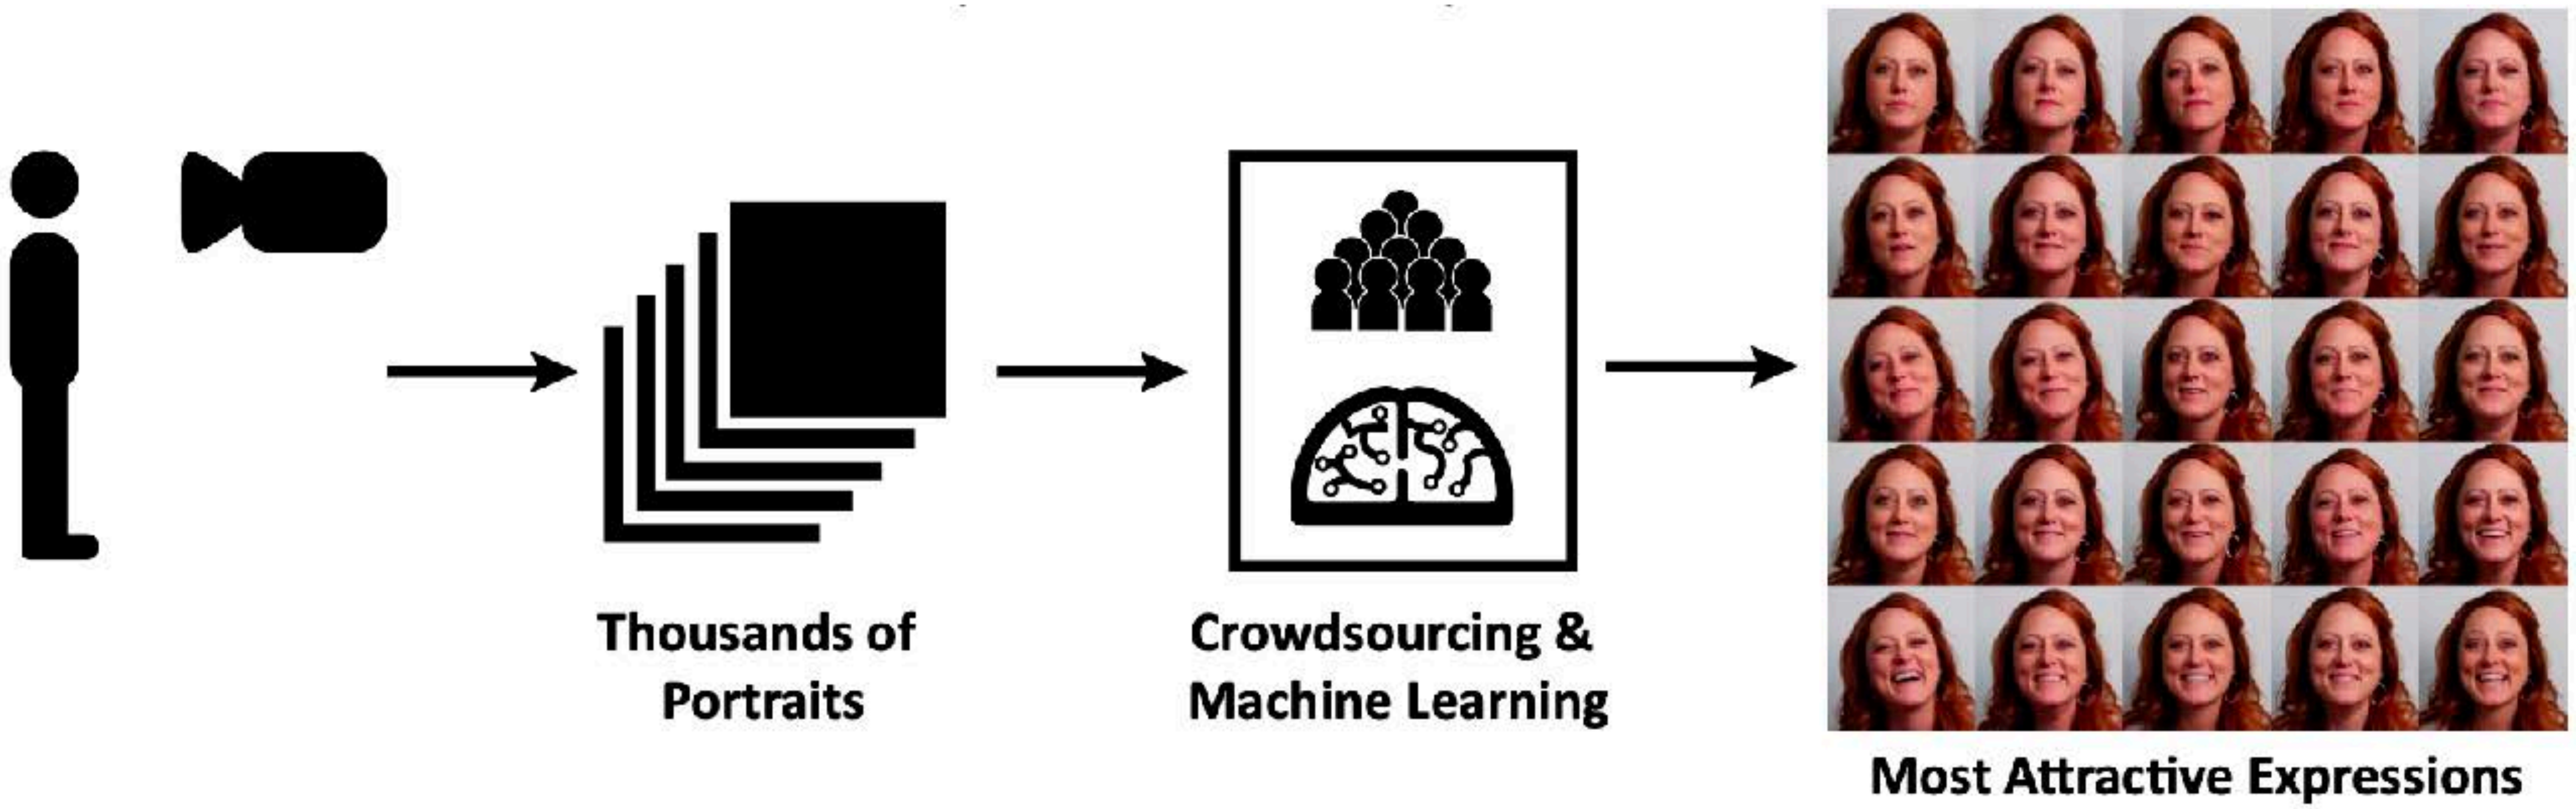

[Zhu+, SA 2014] Jun-Yan Zhu, Aseem Agarwala, Alexei A. Efros, Eli Shechtman, and Jue Wang. 2014. Mirror mirror: crowdsourcing better portraits. ACM Trans. Graph. 33, 6, pp.234:1–234:12 (2014). https://doi.org/10.1145/2661229.2661287

Types of Intelligence [2/4]: Semantic Attributes

- Human-understandable concepts (such as “modern”, “strong”, etc.)
- Based on semantic attributes, users can intuitively explore different visual options

[Yumer+, SIGGRAPH 2015] Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K. Hodgins, and Levent Burak Kara. 2015. Semantic shape editing using deformation handles. ACM Trans. Graph. 34, 4, pp.86:1–86:12 (2015). https://doi.org/10.1145/2766908

Types of Intelligence [3/4]: Perceptual Similarity

E.g., crowdsourcing can be used to measure or estimate perceptual similarity between illustrations, which enables an intelligent search tool

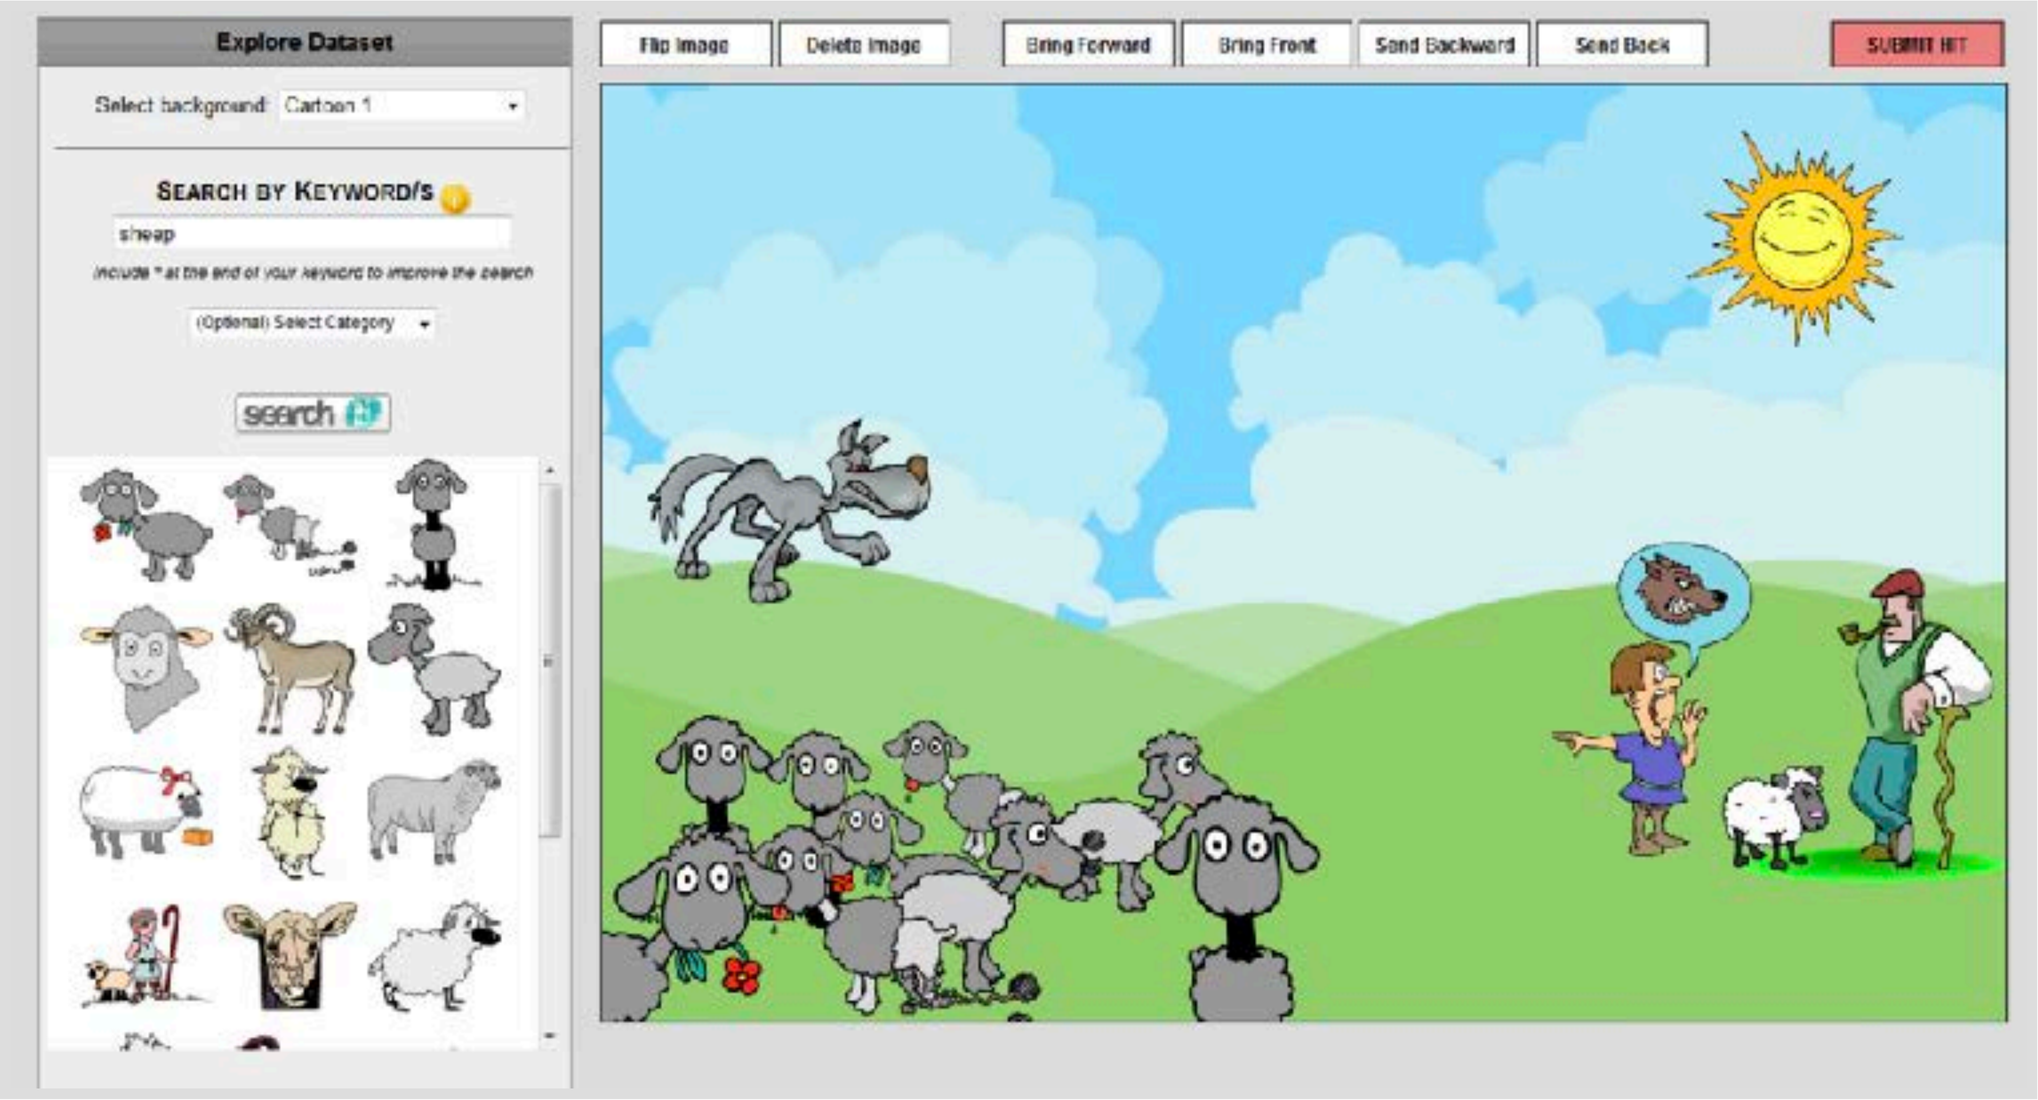

[Garces+, SIGGRAPH 2014] Elena Garces, Aseem Agarwala, Diego Gutierrez, and Aaron Hertzmann. 2014. A similarity measure for illustration style. ACM Trans. Graph. 33, 4, 93:1–93:9 (2014). https://doi.org/10.1145/2601097.2601131

Types of Intelligence [4/4]: Perceptual Compatibility

- Compatibility: how well multiple objects go together
- E.g., by estimating perceptual compatibility between 3D model assets, a scene assembly tool can suggest reasonable variations

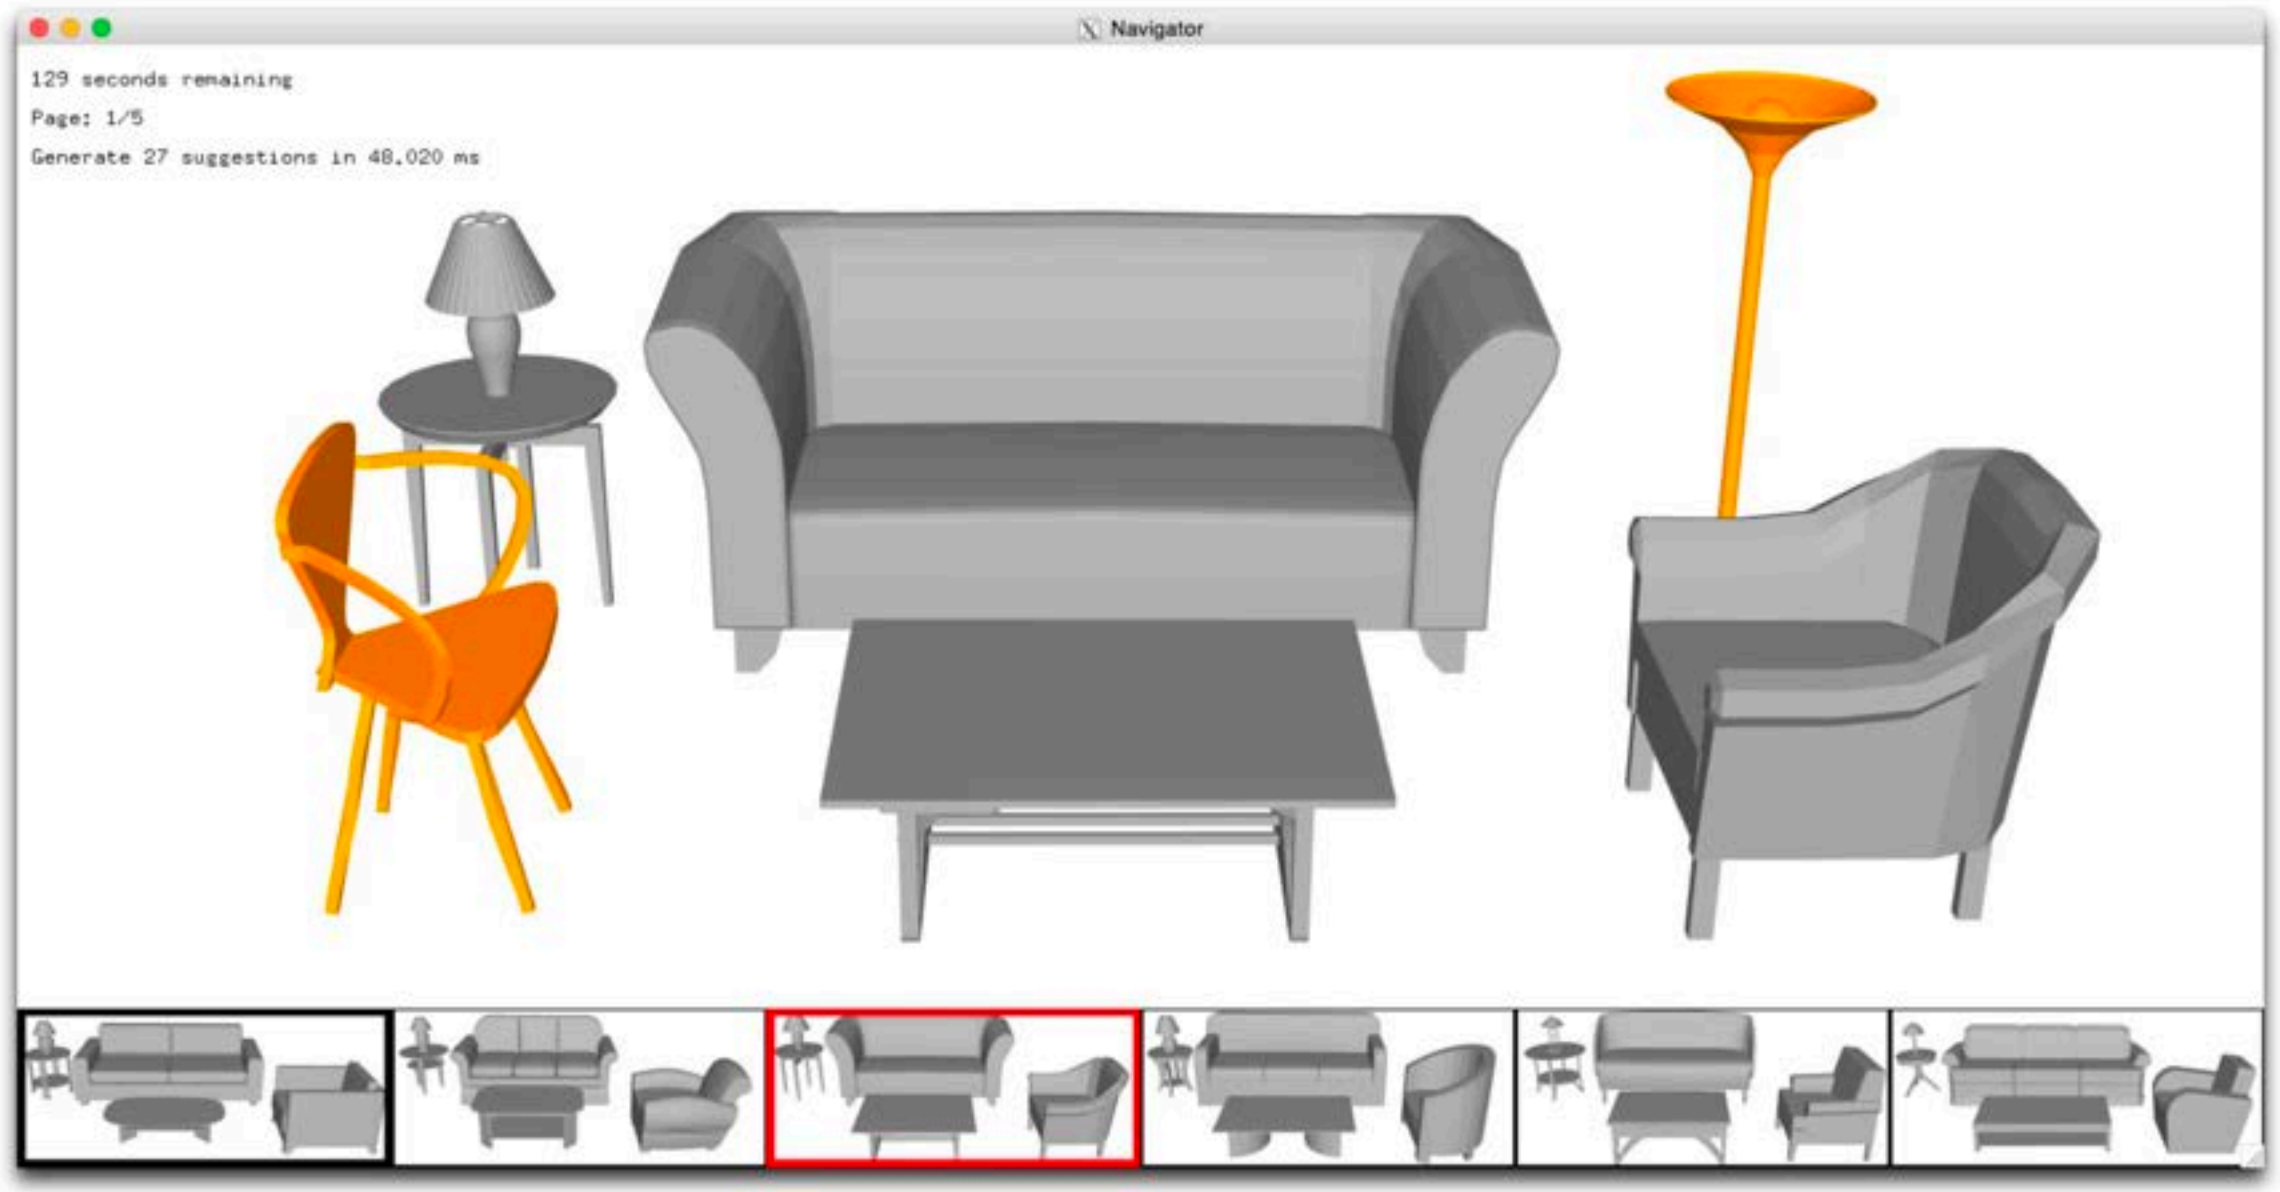

[Liu+, SIGGRAPH 2015] Tianqiang Liu, Aaron Hertzmann, Wilmot Li, and Thomas Funkhouser. 2015. Style compatibility for 3D furniture models. ACM Trans. Graph. 34, 4, pp.85:1–85:9 (2015). https://doi.org/10.1145/2766898

Beyond Parameter Tweaking

Crowd-powered tools can generate **structured perceptual feedbacks** for the on-going design, in an on-demand manner

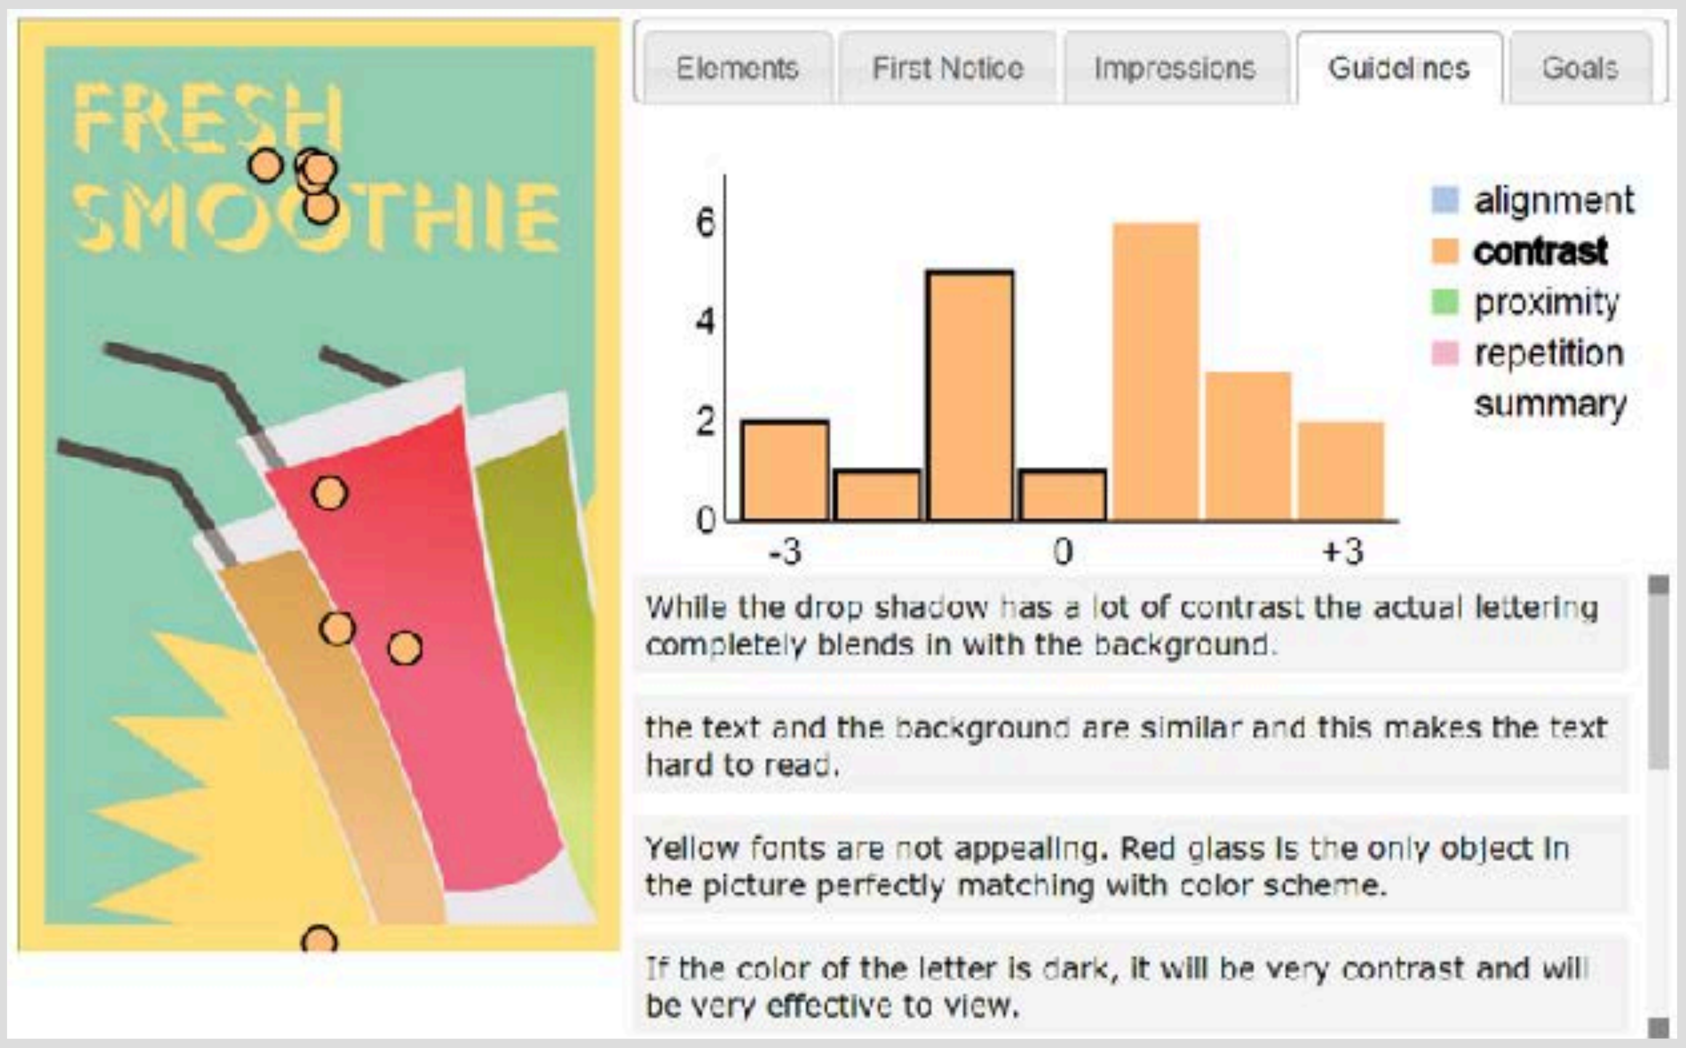

[Xu+, CSCW 2014] Anbang Xu, Shih-Wen Huang, and Brian Bailey. 2014. Voyant: generating structured feedback on visual designs using a crowd of non-experts. In Proc. CSCW '14. pp.1433–1444. https://doi.org/10.1145/2531602.2531604

Another Approach to Incorporate Human Intelligence

Build tools that involve the **user** (instead of **crowds**) in their routine loop

Another Approach to Incorporate Human Intelligence

Build tools that involve the **user** (instead of **crowds**) in their routine loop

- The **crowds-in-the-loop** Bayesian optimization algorithm [Koyama+, SIGGRAPH 2017] was extended for building a new **user-in-the-loop** optimization tool [Koyama+, SIGGRAPH 2020]

[Koyama+, SIGGRAPH 2020] Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential Gallery for Interactive Visual Design Optimization. ACM Trans. Graph. 39, 4, pp.88:1–88:12 (2020). https://doi.org/10.1145/3386569.3392444

Summary

Intelligent Tools for Creative Graphics by Crowdsourcing

Summary

Tools can be more intelligent by crowdsourcing

Summary

Tools can be more intelligent by crowdsourcing

- Crowdsourcing enables tools to **quantify perception-related concepts**
  - E.g., preference, semantic attributes, etc.

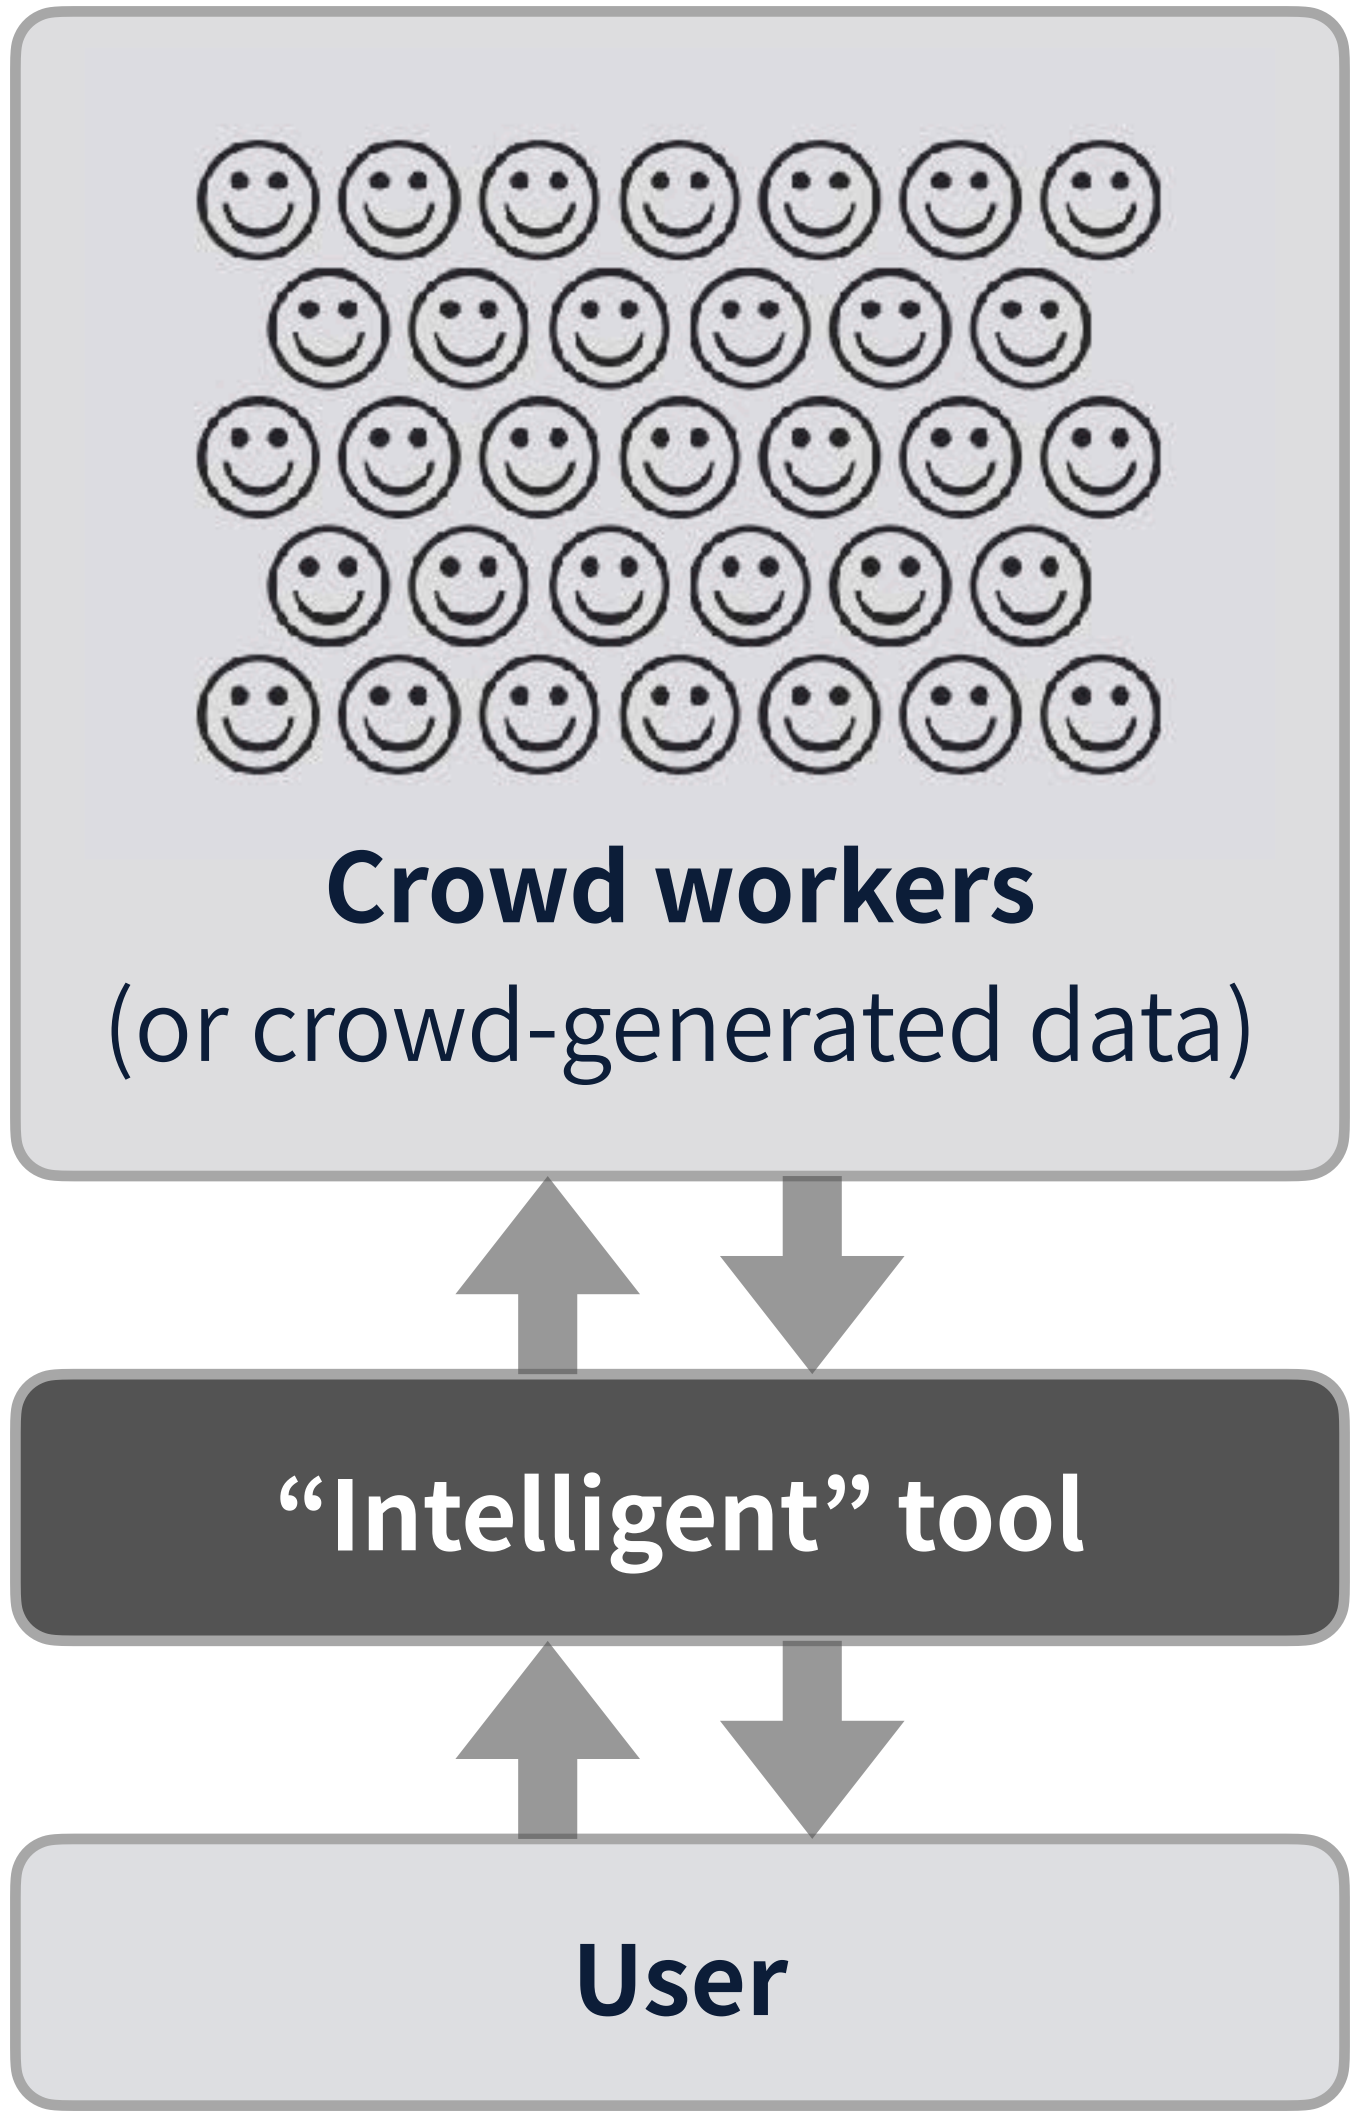

Summary

Tools can be more intelligent by crowdsourcing

- Crowdsourcing enables tools to **quantify perception-related concepts**
  - E.g., preference, semantic attributes, etc.
- Crowd-powered tools can offer functions for **guided exploration**, including …
  - Enhanced sliders with guidance
  - Intelligent suggestions
  - “People’s choice” solver

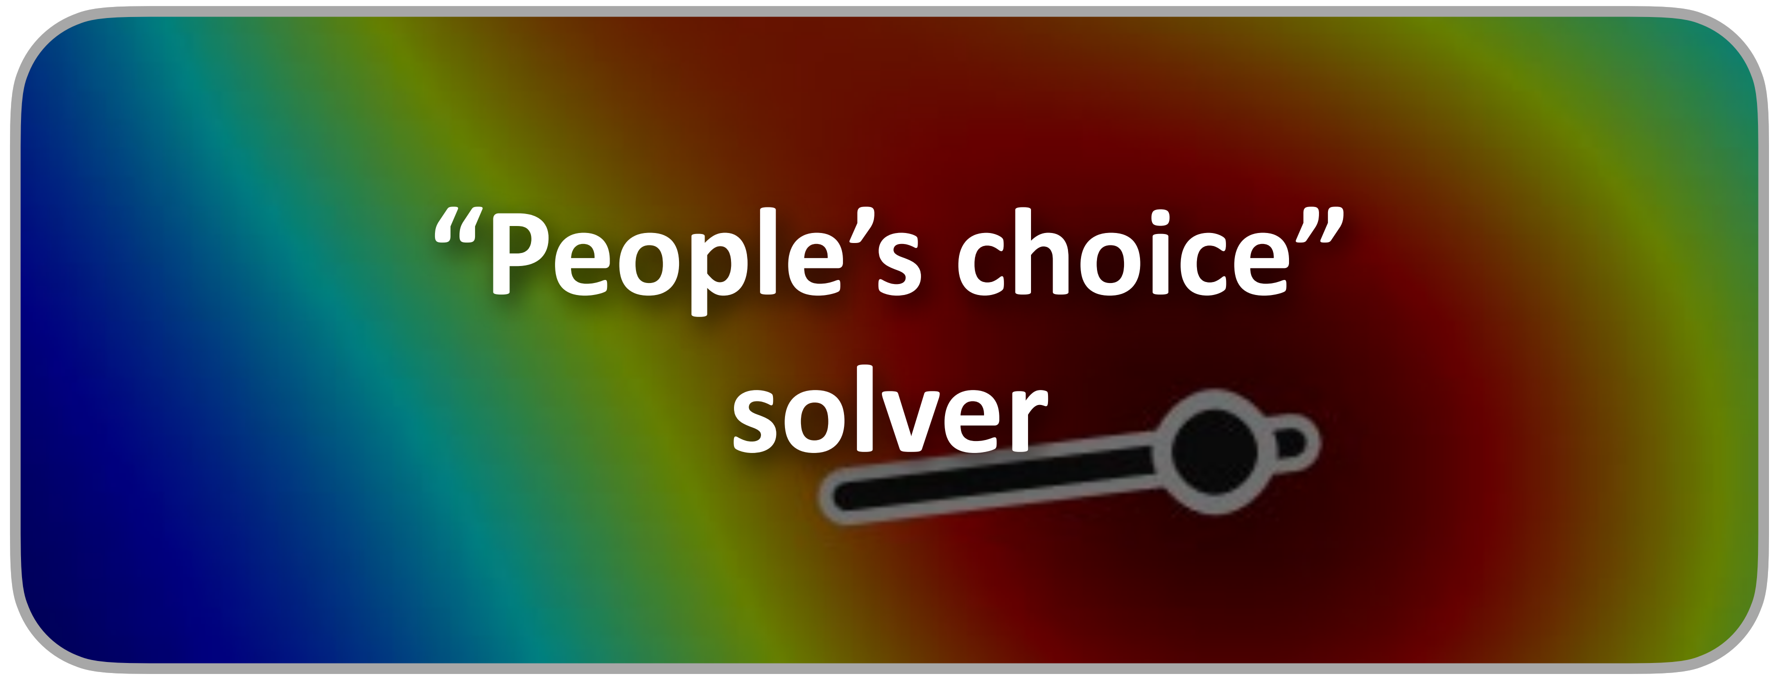

References

- **[Garces+, SIGGRAPH 2014]** Elena Garces, Aseem Agarwala, Diego Gutierrez, and Aaron Hertzmann. 2014. A similarity measure for illustration style. ACM Trans. Graph. 33, 4, 93:1–93:9 (2014). https://doi.org/10.1145/2601097.2601131
- **[Gingold+, TOG (2012)]** Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. 2012. Micro perceptual human computation for visual tasks. ACM Trans. Graph. 31, 5, 119:1–119:12 (2012). https://doi.org/10.1145/2231816.2231817
- **[Koyama+, UIST 2014]** Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. 2014. Crowd-powered parameter analysis for visual design exploration. In Proc. UIST '14. pp.65–74. https://doi.org/10.1145/2642918.2647386
- **[Koyama+, SIGGRAPH 2017]** Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. ACM Trans. Graph. 36, 4, pp.48:1–48:11 (2017). https://doi.org/10.1145/3072959.3073598
- **[Koyama+, Computational Interaction (2018)]** Yuki Koyama and Takeo Igarashi. 2018. Computational Design with Crowds. In Computational Interaction (Eds. Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes), Oxford University Press, pp.153–184. https://arxiv.org/abs/2002.08657
- **[Koyama+, SIGGRAPH 2020]** Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential Gallery for Interactive Visual Design Optimization. ACM Trans. Graph. 39, 4, pp.88:1–88:12 (2020). https://doi.org/10.1145/3386569.3392444
- **[Liu+, SIGGRAPH 2015]** Tianqiang Liu, Aaron Hertzmann, Wilmot Li, and Thomas Funkhouser. 2015. Style compatibility for 3D furniture models. ACM Trans. Graph. 34, 4, pp.85:1–85:9 (2015). https://doi.org/10.1145/2766898
- **[von Ahn, 2005]** Luis von Ahn. 2005. Human Computation. PhD thesis, Carnegie Mellon University, CMU-CS-05-193. http://reports-archive.adm.cs.cmu.edu/anon/2005/CMU-CS-05-193.pdf
- **[Xu+, CSCW 2014]** Anbang Xu, Shih-Wen Huang, and Brian Bailey. 2014. Voyant: generating structured feedback on visual designs using a crowd of non-experts. In Proc. CSCW ’14. pp.1433–1444. https://doi.org/10.1145/2531602.2531604
- **[Yumer+, SIGGRAPH 2015]** Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K. Hodgins, and Levent Burak Kara. 2015. Semantic shape editing using deformation handles. ACM Trans. Graph. 34, 4, pp.86:1–86:12 (2015). https://doi.org/10.1145/2766908
- **[Zhu+, SA 2014]** Jun-Yan Zhu, Aseem Agarwala, Alexei A. Efros, Eli Shechtman, and Jue Wang. 2014. Mirror mirror: crowdsourcing better portraits. ACM Trans. Graph. 33, 6, pp.234:1–234:12 (2014). https://doi.org/10.1145/2661229.2661287

Intelligent Tools for Creative Graphics by Crowdsourcing

Yuki Koyama

National Institute of Industrial Science and Technology (AIST), Japan