D: A Cool Trick for Next-generation Computational Imaging

Abstract

The abstract is a summary of the paper and the main findings. Briefly outline what you are doing, what the specific problem is you are trying to solve, how your approach is different from previous approaches, what the key idea of your approach is, and what difference your approach makes for the outlined problem.

Keywords: computational imaging, computational displays

1 Introduction

Set context Throughout the last few years, technology X has emerged and it is starting to make a significant impact. Here are a whole bunch of applications for technology X to make sure the reader understands just how important this technology is. Now mention what is at the core of this technology (of course that's the small part of the bigger thing that you're actually working on; this will provide the transition to the next paragraph). *Goal of this paragraph: start broad and big, then narrow it down to your specific topic. Heilmeier questions that should be answered: what are you doing? why is it important?*

Problem statement The current generation of technology X works like this (explain the fundamental working principle of technology X's core) [Campbell and G 1960]. Although this is a really successful principle, it also has severe shortcomings in many applications [Fincham 1951; Toates 1972]. Now emphasize the problem and make it absolutely clear, use a few references here. This problem has been address in the past by approaches that can be classified as doing A [Sweeney et al. 2014; Tsuetaki and Schor 1987; Kooi and Toet 2004], B [Heron et al. 2001; Waldkirch et al. 2004; Melzer 1998], or C [Sugihara and Miyasato 1998; Dolgoff 1997]. Unfortunately, due to some reason that should be described here, none of these approaches really solve the problem. *Goal of this paragraph (can also split into two): identify the problem in an absolutely clear way (if the reader misses this, they won't get the point of your paper), give a quick overview of how people have addressed it. Heilmeier question that should be answered: what is the problem with current practice?*

Proposed approach In this paper, we propose a new approach, called D. Based on a really clever insight and on some other technology components that just became available now (so prior work couldn't have possibly done this), we introduce a new approach, to address the core problem of technology X. Figure **??** shows you in detail what the device is that we built, what the best result is that we got, and it also shows you the magic behind the trick. If it is still not clear how it works, we illustrate how it works again in Figure 1, our clever insight is really easy to understand – just look at the figure. Yes, you may have heard of this approach before, because it is commonly used in a few other fields, like medical imaging, time travel, and rocket science. However, we are the first to make the connection to technology X and thoroughly evaluate it in that context. The benefits of using D for X are that we can do things no other approach was able to do, and that we are better at whatever approaches A, B, and C do (see Fig. 2) (*don't get carried away here, stay calm and make sure all arguments have appropriate references cited and are really true - keep in mind that your reviewers are most likely the authors of these papers, so don't insult their work!*). *Goal of this paragraph: give really intuitive insight in your approach, give an overview of why it's better than or different from other approaches and all the cool things you can do with that. Heilmeier questions that should be answered: how does your approach work? how does it compare to current practice?*

Impact Now talk more about the potential impact of your approach and also address how this could be useful for other, non-obvious things. Try to draw connections to other areas in graphics and interactive techniques. Make it relevant for the audience (i.e. specific reviewers). *Goal of this paragraph: help reader understand why they should care and why they should accepted it to SIGGRAPH now (and not later). Heilmeier questions that should be answered: what difference will it make? why should the reviewer care? also, your paper is probably not perfect, so why should the reviewer accept it now and not tell you to make it better and come back next year?*

Contributions Here's a quick summary of the paper. In particular, we make the following contributions: *Goal of this paragraph: help reviewer identify the contributions*

- We *introduce* approach D as a new way of addressing the problem of technology X.
- We *analyze* this approach and show that it is better, faster, and easier that previous approaches.
- We *build* a prototype device to verify this approach and *demonstrate* its practical benefits.
- We *evaluate* approach D with a large-scale user study.

1.1 Overview of Limitations

The reviewer may have a few doubts already: "but wouldn't it fail in situation ..."? Do not wait until the end of the paper to address obvious concerns, because the reviewers will have those in the back of their heads while they are reading through the paper. Best address obvious concerns early on to free the reviewer's mind and let them enjoy the rest of the paper. This doesn't have to be long and is not necessary for all papers, but can be very helpful. *Goal of this paragraph: address possible concerns early on*

2 Related Work

Since the invention of the stereoscope in the 1830s [Wheatstone 1838], significant progress has been made towards making technology X digital and interactive [Sutherland 1968] (cite some more). With the rise of cellphone technology, high-resolution and inexpensive microdisplays have become available as have low-latency inertial measurements units. Together, these developments have led to technology X being practical now. *Goal of this paragraph: give a bit of a historical context of how the field has evolved to what it is today.*

Yet, engineering challenges remain. In particular, technology X requires a vast amount of data that needs to be processed and power consumption along with large device form factors prevent are prohibitive for wearable computing applications. Whereas the latter properties are constantly being improved by the industry, data processing has been an active area of research in the academic community. Generally, approaches to tackle the problem can be classified into approaches A, B, and C. *Goal of this paragraph: tell reader*

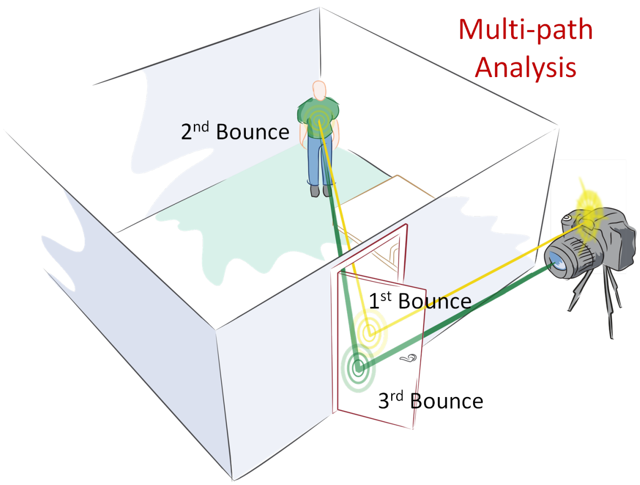

Figure 1: *Here is an intuitive explanation of how the proposed approach works.*

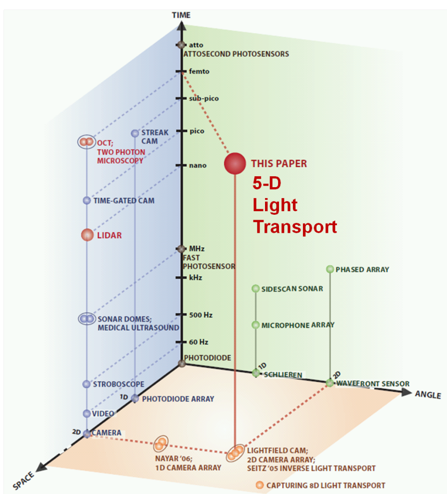

Figure 2: *Here is how the proposed approach compares to previews work and why it is unique/better/faster/easier.*

that there are many problems, but why you picked the one you did. Quick overview of categories of approaches.

Approach A is successful because of many reasons. For example, Sweeney et al. [2014] were the first to make a really cool insight and Tsuetaki and Schor [1987] verified Sweeney's theory with experimental proof. The last approach along these lines was recently described by Kooi et al. [2004] and that's the state of the art for category A. The approach proposed here builds on ideas related to category A, but because all of these approaches do not consider a special case scenario that is really important for technology X, we follow a different school of thought. *Goal of this paragraph: summarize category A with all important papers. Tell the reader how your approach is similar to but different from this category and why that's a better way of doing it.*

We are not the first to consider this special case. [Heron et al. 2001; Waldkirch et al. 2004; Melzer 1998] have all described approaches that use a similar insight. (describe them in more detail) However, they missed something important. We make the insight that a critical piece of the algorithm is missing and demonstrate that this missing piece – D – is actually commonly used in medical imaging, time travel, and rocket science. Here, we show how to apply D to technology X and how it can significantly improve it in many ways. *Goal of this paragraph: describe the set of approaches closest to yours in your area. But then make it clear how your approach is different, potentially borrowing ideas from other fields.*

Approach D is commonly used in medical imaging like computerized tomography, magnetic resonance imaging, and ultrasound [Marran and Schor 1997] (cite a textbook or some comprehensive reference here). In addition, time travel and rocket science also use these ideas in a very specific way that should be described here. Overall, D has proven to be a robust approach to solving the problem that it solves and our goal is to bring these insights to technology x.

This section should be full of references. It's really important to categorize them and discuss these categories in detail. It's also important to identify the one paper or category of approaches that is closest to yours and describe that in detail. You always want to draw a connection between what you do and what each category does at the end of each paragraph so provide smooth transitions to the next paragraph. There should be a logical flow between paragraphs.

3 Math / Theory / System

This section introduces the mathematical foundations of D or if it's a systems paper this section introduces the system and all its components or the theory that your D is based on. If you have a math-oriented imaging paper, oftentimes it makes sense to structure this section into 2-3 parts: image formation, inverse method, and something else.

Always great to include an illustration here that makes it easy for the reader to understand what all the mathematical symbols are. Or you could go into more detail for individual parts of your system with a little figure. If you have some algorithm, include pseudo code to make it more reproducible!

3.1 Image Formation

Here, we will start with some equation that is intuitive for the reader and then build on that.

The resolution of diffraction-limited imaging systems is governed by Abbe's diffraction limit [Born and wolf 1997]. Light of wavelength λ, traveling in a medium with refractive index n and converging to a spot with angle θ will make a spot with radius

$$d = \frac{\lambda}{2n\sin\theta}. \tag{1}$$

The denominator $n \sin\theta$ is called the numerical aperture (NA). The spread of the diffraction-limited point spread function (PSF) ρ is approximated by the diameter of the first null of the Airy disk $d/2$. An captured image b is then modeled by the convolution of the PSF and the object o. In fluorescence microscopy, the illumination pattern s exciting the fluorescent object must be taken into account, resulting in

$$b = \rho * (o \cdot s), \tag{2}$$

where $*$ represents a convolution and $\cdot$ is the Hadamard product (i.e., element-wise multiplication).

In the following, we consider an optical apparatus where multiple images b_i are recorded of the same object but changing illumination patterns. In particular, we focus on speckle-based illumination patterns that are created by a diascopic (transmitted) where a coherent source illuminates a thin diffuser that is mounted at a small distance below the object (see Figure 3). In successively captured images $i = 1 \ldots N$, the light source is laterally shifted by a small amount. Due to the memory effort [], the resulting speckle pattern on the object is a shifted. We denote the deformation between a reference speckle pattern s of the light source in the center and some shifted version of it as the transport operator $T_i(s)$, such that image i is given as

$$b_i = \rho * (o \cdot T_i(s)). \tag{3}$$

3.2 Inverse Methods

Now we either discretize the image formation (or we already did that at the end of the last subsection), and we talk about how we develop and objective function around this discrete model. Then we propose some way of solving the objective.

The goal is now to estimate both the high-resolution object and the unknown speckle illumination pattern from a sequence of measurements. For this purpose, we discretize all variables and form the following objective function

$$\begin{array}{ll} \underset{\{\mathbf{o},\mathbf{s}\}}{\text{minimize}} & J(\mathbf{o},\mathbf{s}) = \sum_{i=1}^{N} \frac{1}{2} \left\| \mathbf{b}_i - \mathbf{P}\left(\mathbf{o} \cdot T_i(\mathbf{s})\right) \right\|_2^2 \\ \text{subject to} & 0 \leq \mathbf{o},\mathbf{s} \end{array} \tag{4}$$

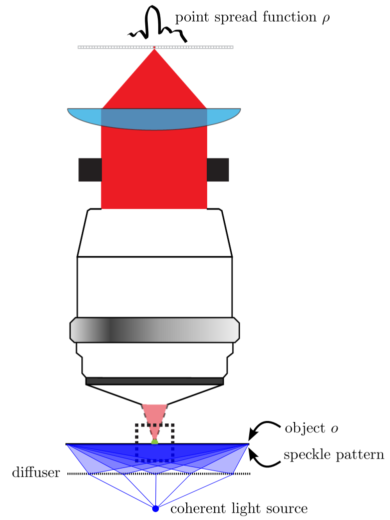

Figure 3: *Here's an intuitive illustration of all quantities used in the equations.*

Here, $\mathbf{b}_i, \mathbf{o}, \mathbf{s} \in \mathbb{R}_+^M$ where M is the number of sensor pixels. Without loss of generality, we write the discrete convolution with ρ as a product with matrix $\mathbf{P} \in \mathbb{R}^{M \times M}$. To solve the objective, we derive the gradient of $J(\mathbf{o},\mathbf{s})$ with respect to $\mathbf{o}$ and $\mathbf{s}$ as

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{o}} &= \left(\sum_{i=1}^{N} \mathbf{S}_i^T \mathbf{P}^T \mathbf{P} \mathbf{S}_i \right) \mathbf{o} - \sum_{i=1}^{N} \mathbf{S}_i^T \mathbf{P}^T \mathbf{b}_i \\ \frac{\partial J}{\partial \mathbf{s}} &= \left(\sum_{i=1}^{N} \mathbf{T}_i^T \mathbf{O}^T \mathbf{P}^T \mathbf{P} \mathbf{O} \mathbf{T}_i \right) \mathbf{s} - \sum_{i=1}^{N} \mathbf{T}_i^T \mathbf{O}^T \mathbf{P}^T \mathbf{b}_i. \end{aligned} \tag{5}$$

In this formulation, the diagonal matrices $\mathbf{O} \in \mathbb{R}^{M \times M}$ and $\mathbf{S}_i \in \mathbb{R}^{M \times M}$ contain $\mathbf{o}$ and $T_i(\mathbf{s})$ on their diagonals, respectively. The matrix $\mathbf{T}_i \in \mathbb{R}^{M \times M}$ encodes the transform T_i. Although Equation 4 is not convex, it is bi-convex, i.e., by fixing either $\mathbf{o}$ or $\mathbf{s}$, the formulation is convex with respect to the other variable. Therefore, we can employ an alternating least-squares scheme where we initialize $\mathbf{o}$ and $\mathbf{s}$ with some random guess and then iteratively updating them in an alternating fashion. This update scheme is shown in Algorithm 1 and each of the individual steps can be conveniently solved with any non-negative least squares solver; we use the sparse, constrained, large-scale trust region method implemented in MATLAB's *lsqlin* function [Coleman and Li 1996].

Algorithm 1 Alternating Least Squares

```
1: function [o, s] = ALS(o, s)
2:   for each iteration
3:     o ← arg min_{0<o} (Σ_{i=1}^N S_i^T P^T P S_i) o − Σ_{i=1}^N S_i^T P^T b_i
4:
5:     s ← arg min_{0<s} (Σ_{i=1}^N T_i^T O^T P^T P O T_i) s − Σ_{i=1}^N T_i^T O^T P^T b_i
6:
7:   end
8: end
```

3.3 Bounds on Resolution

If this is a general imaging paper, we now know how to recover the info we want from our measurements. So it may be good to either *derive some fundamental performance bounds here or talk about optimizing the optical setup for this algorithm.*

Here, we derive a bound on the fundamental limits of resolution improvement of speckle-based coded illumination fluorescence microscopy. It is well known that structured illumination microscopy (SIM) can improve the resolution of an object by a factor of $2\times$ compared to the diffraction limit [Gustafsson 2000]. This limit, however, is based on the assumption that that the condenser lens of the illumination has the same NA as the detection objective. While this is assumption makes sense for many applications, we advocate for random speckle illumination. There is something fundamentally different in the image formation here that changes the

$$d = \frac{\lambda}{2n\phi\left(\xi\right)\sin\theta}. \tag{6}$$

We see that the resolution in the proposed system is controlled by some optical parameter that is expressed as $\phi\left(\xi\right)$. In theory, the resolution of fluorescence samples could be improved by a factor of $100\times$ over the diffraction limit while providing a field of view comparable to that of the detection objective. We analyze these benefits in the following sections with simulations and a prototype device.

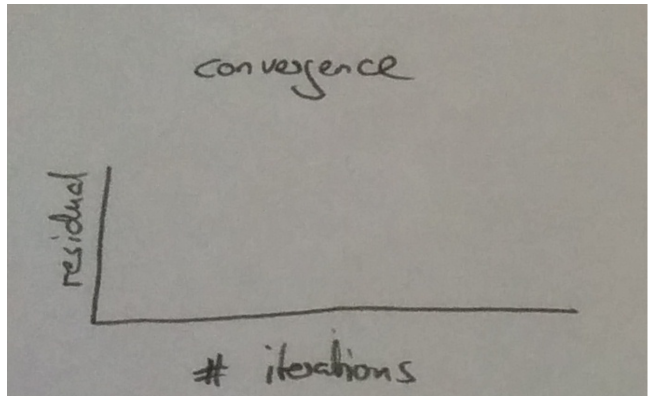

Figure 4: *convergence - should go down.*

4 Implementation

4.1 Software

We implemented everything in MATLAB. For a 25 megapixel image, the algorithm takes 10 secs to converge in 500 iterations; we show convergence in Figure 4. We hand-tuned the parameters to $\lambda_1 = 0.001$ and $\sigma = 12.45$. All source code and data will be made available (*this addresses the reproducibility issue*).

4.2 Hardware

We built a really cool prototype. Now list all the prototype parts and specific model numbers. The hardware should be fully reproducible from the information given here! You can see a photograph of the prototype in Figure 5.

4.3 Calibration

Describe in detail how you calibrated the hardware setup!

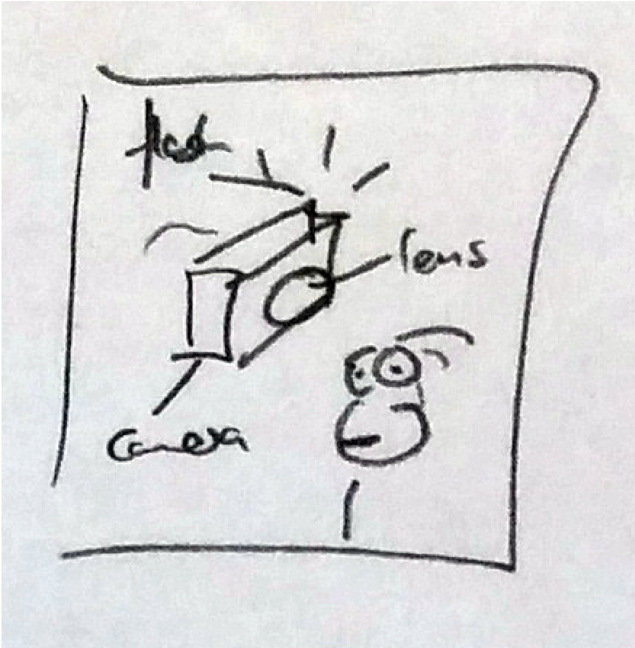

Figure 5: *Show photo of prototype and label parts.*

5 Assessment

5.1 Analysis

Evaluate all algorithmic or hardware parameters here.

5.2 Interpretation

Potentially analysis your decompositions, patterns, or whatever you do and try to make them intuitive. What does this complicated algorithm do intuitively?

5.3 Results

Show visually compelling results here. Make sure you show different results, i.e. if you have multiple scenes make sure that every scene or result tells its own story!

5.4 Comparison to Previous Work

You need both qualitative and quantitative comparisons. The qualitative comparison must be absolutely and unmistabably obvious! If it's hard to see, zoom in and magnify what they are supposed to look at.

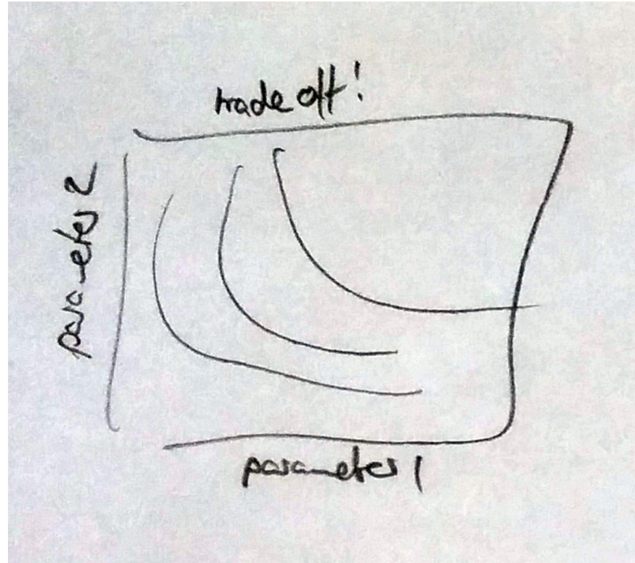

Figure 6: *Show some plot that shows that the effects of the parameters are.*

6 Bells & Whistles

Mandatory video to watch on youtube: shamwow (`https://www.youtube.com/watch?v=QwRISkyV_B8`).

Now that we have focused on introducing method D, deriving a bunch of math, analyzing all the algorithmic parameters, comparing it to previous work, and showing a whole lot of really cool results for the core thing that this paper is focused on, let us show you a bunch of extra stuff that you can also do with the proposed approach.

Remember the shamwow video. This is basically a wiping cloth and for the past couple of minutes in the video clip (or the last 3 sections of the paper), we have focused on the main application: wiping stuff off all possibly conceivable surfaces. But now we'll show you a bunch of other stuff that you get for free! For example, you can use shamwow as a mop (wiping cloth with a stick) so you can clean your entire house, you can also use shamwow for wiping your cat or dog or your bird, you can clean your car, and finally, you can also use it as a diaper and as a blanket.

Be concise in this section, because the paper is about the previous 3 sections. Just write a few paragraphs and show the best results for each bell & whistle.

7 Discussion

In summary, we proposed a new thing, analyzed it in detail, showed a bunch of results, built a prototype, etc etc. Basically write a paragraph summarizing the paper. By now the reader may be asleep and they probably missed a couple of important details when they skimmed over the technical sections, so just briefly summarize everything in one paragraph.

Make sure to highlight all the benefits of your approach in this summary! You can even write a second paragraph just discussing these.

Limitations However, as with all methods, since there would never be any future papers in this area, our method also has limitations. First, our approach makes the following fundamental as-

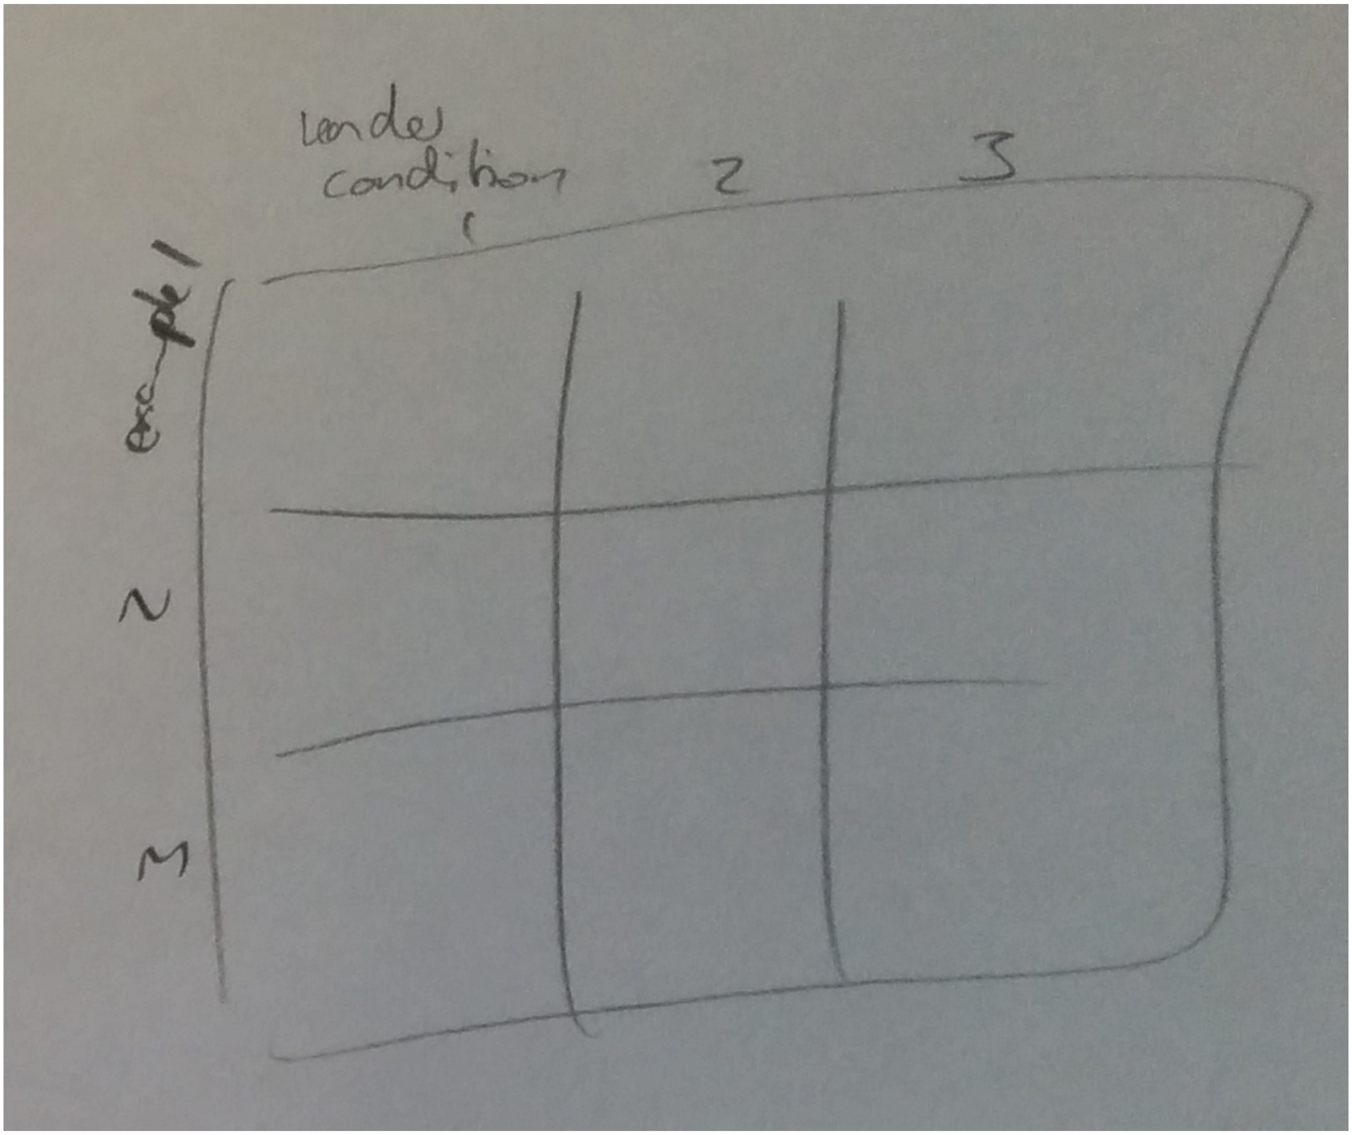

Figure 7: *Show results for multiple examples here. Each one should tell a different story!*

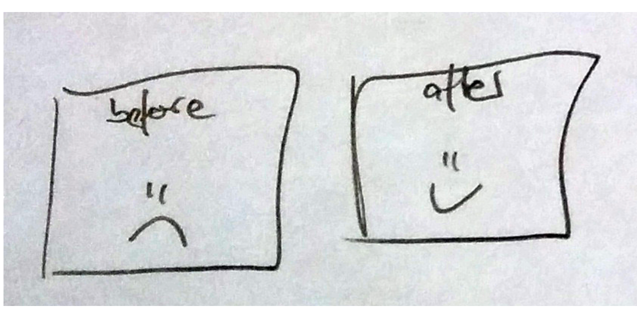

Figure 8: *This has to be absolutely obvious!*

sumptions ... If these assumptions are not valid, as is the case in scenarios x, y, and z, then this method may not work. Second, our methods also fails in the following scenarios: ... An example of *how* it fails is shown in Figure 9.

Now discuss a few possible strategies for overcoming these limitations here (or just integrate them into the previous paragraph so as to dismiss the limitations as easily fixable). But focus on possible methods to fix exactly what's wrong with your current approach.

Future Work In the future, we will actually fix all the things that currently don't work. We would also like to apply the approach proposed here to the completely different problem of ... we believe this has a lot of promise. And we want to do a bunch of other stuff as well.

8 Conclusion

The conclusion is different from the discussion. In the discussion, we talked about what we did and what the benefits and limitations are. We talked about future work and tradeoffs that need to be made. The conclusion is different – write a paragraph that you want people to take away from this paper, placing it into the bigger context. No summary of what you've done or any details. Think about it as this

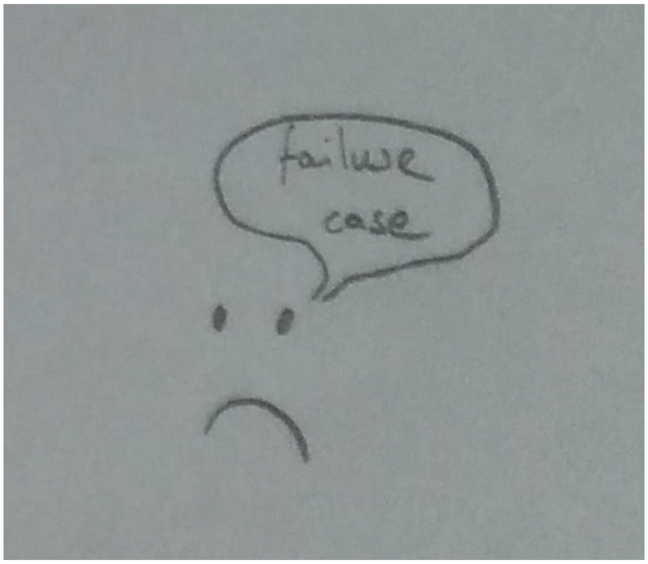

Figure 9: *Failure case. Don't just show that it fails here but* how it fails!

"now that we know all the things discussed in the paper, how will it (potentially) change common practice, how will it (potentially) enable new stuff? What is the bigger goal towards which this paper makes a significant step?". So be concise and end on a big note that people will remember alluding to the bigger vision.

References

BORN, M., AND WOLF, E. 1997. *Principles of Optics*. Cambridge University Press.

CAMPBELL, F. W., AND G, W. 1960. Dynamics of accommodation responses of the human eye. *Journal of physiology 151*, 285–295.

COLEMAN, T. F., AND LI, Y. 1996. A reflective newton method for minimizing a quadratic function subject to bounds on some of the variables. *SIAM Journal on Optimization 6*, 4, 1040–1058.

DOLGOFF, E. 1997. Real-depth imaging: a new 3d imaging technology with inexpensive direct-view (no glasses) video and other applications. *Proc. SPIE 3012*, 282–288.

FINCHAM, E. F. 1951. The accommodation reflex and its stimulus. *British Journal of Ophthalmology 35*, 381–393.

GUSTAFSSON, M. G. 2000. Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy. *J Microscopy 198*, 82–7.

HERON, G., CHARMAN, W., AND SCHOR, C. 2001. Age changes in the interactions between the accommodation and vergence systems. *Optom Vis Sci 10*, 78, 754–62.

KOOI, F. L., AND TOET, A. 2004. Visual comfort of binocular and 3d displays. *Displays 25*, 99 – 108.

MARRAN, L., AND SCHOR, C. 1997. Multiaccommodative stimuli in vr systems: problems and solutions. *Human Factors 39*, 3, 382–388.

MELZER, J. E., 1998. Overcoming the field-of-view/resolution invariant in head-mounted displays.

SUGIHARA, T., AND MIYASATO, T. 1998. 32.4: A lightweight 3-d hmd with accommodative compensation. *SID Digest 29*, 1, 927–930.

SUTHERLAND, I. E. 1968. A head-mounted three dimensional display. In *Proc. Fall Joint Computer Conference*, 757–764.

SWEENEY, L. E., SEIDEL, D., DAY, M., AND GRAY, L. S. 2014. Quantifying interactions between accommodation and vergence in a binocularly normal population. *Vision research 105*, 121–129.

TOATES, F. M. 1972. Accommodation function of the human eye. *Physiological Reviews 52*, 828–863.

TSUETAKI, T., AND SCHOR, C. 1987. Clinical method for measuring adaptation of tonic accommodation and vergence accommodation. *American journal of optometry and physiological optics 64*, 6, 437–449.

WALDKIRCH, M., LUKOWICZ, P., AND TRÖSTER, G. 2004. Multiple imaging technique for extending depth of focus in retinal displays. *Optics Express 12*, 25.

WHEATSTONE, C. 1838. Contributions to the physiology of vision. part the first. on some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London 128*, 371–394.