

From Zero to Hero: Convincing with Extremely Complicated Math

Maximilian Weiherer Bernhard Egger

Funky-Amusing-University (FAU) Erlangen-Nürnberg

maximilian.weiherer@fau.de
bernhard.egger@fau.de

way networks have not demonstrated accuracy gains with extremely increased depth (e.g., over 100 layers).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to be fit by a few stacked layers (not necessarily the entire net), with x denoting the inputs to the first of these layers. If one hypothesizes that multiple nonlinear layers can asymptotically approximate complicated functions¹, then it is equivalent to hypothesize that they can asymptotically approximate the residual functions, i.e., $F(x) = x$ (assuming that the input and output are of the same dimensions). So rather than expect stacked layers to approximate $F(x)$, we explicitly let these layers approximate a residual function $Z(x) := F(x) - x$. The original function thus becomes $F(x) = Z(x) + x$. Although both forms should be able to asymptotically approximate the desired function (an hypothesis), the ease of learning might be different.

This reformulation is motivated by the counterintuitive phenomena about the degradation problem (Fig. 1, left). As we discussed in the introduction, if the added layers can be constructed as identity mappings, a deeper model should have training error no greater than its shallower counterpart. The degradation problem suggests that the solvers might have difficulties in approximating identity mappings by multiple nonlinear layers. With the residual learning reformulation, if identity mappings are optimal, the solvers may simply drive the weights of the multiple nonlinear layers toward zero to approach identity mappings.

In real cases, it is unlikely that identity mappings are optimal, but our reformulation may help to precondition the problem. If the optimal function is closer to an identity mapping than to a zero mapping, it should be easier for the solver to find the perturbations with reference to an identity mapping, than to learn the function as a new one. We show by experiments (Fig. 7) that the learned residual functions in general have small responses, suggesting that identity mappings provide reasonable preconditions.

3.2. Identity Mapping by Shortcuts

We adopt residual learning to every few stacked layers. A building block is shown in Fig. 2. Formally, in this paper we consider a building block defined as:

$$y = F(x, [W], 1) + x. \quad (1)$$

Here x and y are the input and output vectors of the layers considered. The function $F(x, [W], 1)$ represents the residual mapping to be learned. For the example in Fig. 2 that has two layers, $F = [W_1, W_2](x)$ in which \odot denotes

¹This hypothesis, however, is still an open question. See [2].

way networks have not demonstrated accuracy gains with extremely increased depth (e.g., over 100 layers).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to be fit by a few stacked layers (not necessarily the entire net), with x denoting the inputs to the first of these layers. If one hypothesizes that multiple nonlinear layers can asymptotically approximate complicated functions¹, then it is equivalent to hypothesize that they can asymptotically approximate the residual functions, i.e., $F(x) = x$ (assuming that the input and output are of the same dimensions). So rather than expect stacked layers to approximate $F(x)$, we explicitly let these layers approximate a residual function $Z(x) := F(x) - x$. The original function thus becomes $F(x) = Z(x) + x$. Although both forms should be able to asymptotically approximate the desired function (an hypothesis), the ease of learning might be different.

This reformulation is motivated by the counterintuitive phenomena about the degradation problem (Fig. 1, left). As we discussed in the introduction, if the added layers can be constructed as identity mappings, a deeper model should have training error no greater than its shallower counterpart. The degradation problem suggests that the solvers might have difficulties in approximating identity mappings by multiple nonlinear layers. With the residual learning reformulation, if identity mappings are optimal, the solvers may simply drive the weights of the multiple nonlinear layers toward zero to approach identity mappings.

In real cases, it is unlikely that identity mappings are optimal, but our reformulation may help to precondition the problem. If the optimal function is closer to an identity mapping than to a zero mapping, it should be easier for the solver to find the perturbations with reference to an identity mapping, than to learn the function as a new one. We show by experiments (Fig. 7) that the learned residual functions in general have small responses, suggesting that identity mappings provide reasonable preconditions.

3.2. Identity Mapping by Shortcuts

We adopt residual learning to every few stacked layers. A building block is shown in Fig. 2. Formally, in this paper we consider a building block defined as:

$$y = \sum_{i=1}^n \left(\prod_{j=1}^i \text{tanh}(x, W_j) \cdot \left(\frac{\log(1 + \frac{1}{\|W_j\|})}{1 + \exp(-x, W_j)} \right) \right) + x. \quad (2)$$

Here x and y are the input and output vectors of the layers considered. The function $F(x, [W], 1)$ represents the

¹This hypothesis, however, is still an open question. See [2].

Zero success

Hero success

We present zero2hero, an innovative system that turns every scientific paper into an award-winning masterpiece. Given the fact that papers solely using notoriously simple math provably lead to failure (top-tier conference rejections and rude reviewers, diminished respect and appreciation from almost everyone, decline in social status, etc.), zero2hero reliably over-complicates equations so that no one, including yourself, is able to understand what's happening or what ever happened. Buckle up [9] and let zero2hero boost your career, now.

Abstract

Becoming a (super) hero is almost every kid's dream. During their sheltered childhood, they do *whatever it takes* to grow up to be one. Work hard, play hard – all day long. But as they're getting older, distractions are more and more likely to occur. They're getting off track. They start discovering what is feared as *simple math*. Finally, they end up as a researcher, writing boring, non-impressive papers all day long because they only rely on simple mathematics. No top-tier conferences, no respect, no groupies. Life's over.

To finally put an end to this tragedy, we propose a fundamentally new algorithm, dubbed zero2hero, that turns every research paper into a scientific masterpiece. Given a \LaTeX document containing ridiculously simple math, based on next-generation large language models, our system *automatically over-complicates every single equation* so that no one, including yourself, is able to understand what the hell is going on. Future reviewers will be blown away by the complexity of your equations, immediately leading to acceptance. zero2hero gets you back on track, because *you deserve to be a hero*TM. Code leaked at <https://github.com/mwe/ierer/zero2hero>.

1. Introduction

Simple math doesn't impress anybody, neither your grandma nor any reviewer. Scientific papers overgrown with ridiculously underwhelming mathematics are deadly boring to read, often dismissed as trivial, and, ultimately, don't cause the urgently needed pain most readers are desperately looking for. Authors of those 'research' papers also frequently complain about not being treated with the necessary respect, which often manifests itself in the fact that simply too many scientists can follow their 'ideas', or, even worse, are able to suggest improvements to the author's 'work'. How dare they!

As if that weren't enough, it recently has been proven (the proof is left as an exercise for the reader) that getting into top-tier conferences like CVRP, ICCV/ECCV, and NeurISP, depends *solely* on the complexity of the used mathematics, simply judged by counting equations xor inspecting notation. As a consequence, authors who wish to publish at those conferences began to maximize the number of equations and the notational complexity in order to satisfy the reviewer's fetish. Popular tricks to make a paper's math look more complex include, for instance, maximizing the occurrences of Greek letters (try to use as many of them as possible already in

We have here an annotated academic paper with correct source code (100% hit rate).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to y . In this case, we can think of the network as a composition of multiple layers, each of which can be approximated by a linear transformation. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem.

3.2. Identity Mapping by Shortcut

We apply residual learning to the actual ResNet architecture. We use a residual block with two parallel paths. The first path is a convolutional layer with kernel size k and stride s . The second path is a convolutional layer with kernel size k and stride s . We then add the two paths together. This is the residual block. We can then write the overall function as a composition of these blocks, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these blocks approximates the target function $f(x)$. This is the approximation problem.

We have here an annotated academic paper with correct source code (100% hit rate).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to y . In this case, we can think of the network as a composition of multiple layers, each of which can be approximated by a linear transformation. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem.

3.2. Identity Mapping by Shortcut

We apply residual learning to the actual ResNet architecture. We use a residual block with two parallel paths. The first path is a convolutional layer with kernel size k and stride s . The second path is a convolutional layer with kernel size k and stride s . We then add the two paths together. This is the residual block. We can then write the overall function as a composition of these blocks, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these blocks approximates the target function $f(x)$. This is the approximation problem.

We have here an annotated academic paper with correct source code (100% hit rate).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to y . In this case, we can think of the network as a composition of multiple layers, each of which can be approximated by a linear transformation. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem.

3.2. Identity Mapping by Shortcut

We apply residual learning to the actual ResNet architecture. We use a residual block with two parallel paths. The first path is a convolutional layer with kernel size k and stride s . The second path is a convolutional layer with kernel size k and stride s . We then add the two paths together. This is the residual block. We can then write the overall function as a composition of these blocks, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these blocks approximates the target function $f(x)$. This is the approximation problem.

We have here an annotated academic paper with correct source code (100% hit rate).

3. Deep Residual Learning

3.1. Residual Learning

Let us consider $f(x)$ as an underlying mapping to y . In this case, we can think of the network as a composition of multiple layers, each of which can be approximated by a linear transformation. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem. We can then write the overall function as a composition of these layers, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these layers approximates the target function $f(x)$. This is the approximation problem.

3.2. Identity Mapping by Shortcut

We apply residual learning to the actual ResNet architecture. We use a residual block with two parallel paths. The first path is a convolutional layer with kernel size k and stride s . The second path is a convolutional layer with kernel size k and stride s . We then add the two paths together. This is the residual block. We can then write the overall function as a composition of these blocks, $f(x) = h_n \circ h_{n-1} \circ \dots \circ h_1(x)$. The approximation problem is to find a set of parameters θ such that the composition of these blocks approximates the target function $f(x)$. This is the approximation problem.

We have here an annotated academic paper with correct source code (100% hit rate).

4. SEAM: A Simple and Efficient Architecture for Pseudo-Label Generation

2.1. Semantic Pseudo-Label Generation

Our pseudo-label generation process consists of two separate components. We first describe the procedure behind training the 'objectness' network which is designed to obtain detailed boundary information for any object-like region. Next, we describe two different techniques for generating semantic pseudo-labels by combining the output of the objectness network with semantic region proposals, which are obtained from other image-level class labels or bounding box annotations.

Training an Objectness Network

Pixel objectness [24] quantifies how likely a pixel belongs to an object of the class c , other than 'stuff' (class like background, grass, sky, sidewalk, etc.), and should be high over the object unseen during training. We use DeepLabV3+ [23] as a source dataset, \mathcal{D}_c , to train an objectness prior for the 'thing' label. We use a weak form of the COCOStuff dataset, denoted as COCO-Binary and consider it as the source dataset, \mathcal{D}_s . More specifically, we generate COCO-Binary by removing all the three categories assigned to the label one, and the rest categories to zero. We then train the objectness network, θ_c , on the source dataset, \mathcal{D}_s , under two different settings which compare a pixel-wise 'objectness score' (similar to the saliency detection model). In the first setting, we include all the images from the source data, \mathcal{D}_s , regardless of whether the objects found in \mathcal{D}_s images overlap with target data, \mathcal{D}_c . In the second setting, we create a subset of \mathcal{D}_s by excluding those images containing any categories which overlap with \mathcal{D}_c categories. We can formalize the overlapping and non-overlapping settings as follows:

$$O = \begin{cases} \mathcal{D}_s & \text{overlapping} \\ \mathcal{D}_s \setminus \mathcal{D}_c & \text{non-overlapping} \end{cases}$$

where \mathcal{D}_c denotes the set of object classes contained in COCO-Binary used to train the objectness model, θ_c . \mathcal{D}_s represents the non-overlapping subset where there is no semantic category overlap between \mathcal{D}_s and \mathcal{D}_c . Note that the semantic annotations are only used to generate the subset of non-overlapping data, $\mathcal{D}_s \setminus \mathcal{D}_c$, and is not required for training the objectness model, θ_c . We believe the non-overlapping setting is more challenging than image-level based WSS [55, 76, 79], because these methods contain semantic overlap within the source and target data. In both settings, we train the objectness classifier using the cross-entropy approximation and use the binary cross-entropy loss function. The main goal of the objectness classifier is to learn a strong objectness prior for the 'thing' label.

Class-Driven Pseudo-Labels

CAM [83] is widely used as a weak source of supervision as it roughly localizes semantic object areas. Following previous works [13], we first generate CAMs for training images by adapting the method of [83] using a multi-label image classification network. For a fair comparison, we use a ResNe50 [24] model as the classification network, we used in other CAM-based methods [2, 3, 76]. We directly utilize these CAMs to generate pseudo-labels by thresholding their confidence scores for each class label at every pixel predicted to be an object by the class-agnostic objectness network (see Fig. 1(b)). We can formalize this procedure as follows:

$$p_c = \begin{cases} \max_{i \in \{1, \dots, C\}} \text{CAM}_i(x, y) & \text{if } \text{CAM}_c(x, y) > \tau \\ 0 & \text{otherwise} \end{cases}$$

We have here an annotated academic paper with correct source code (100% hit rate).

4. SEAM: A Simple and Efficient Architecture for Pseudo-Label Generation

2.1. Semantic Pseudo-Label Generation

Our pseudo-label generation process consists of two separate components. We first describe the procedure behind training the 'objectness' network which is designed to obtain detailed boundary information for any object-like region. Next, we describe two different techniques for generating semantic pseudo-labels by combining the output of the objectness network with semantic region proposals, which are obtained from other image-level class labels or bounding box annotations.

Training an Objectness Network

Pixel objectness [24] quantifies how likely a pixel belongs to an object of the class c , other than 'stuff' (class like background, grass, sky, sidewalk, etc.), and should be high over the object unseen during training. We use DeepLabV3+ [23] as a source dataset, \mathcal{D}_c , to train an objectness prior for the 'thing' label. We use a weak form of the COCOStuff dataset, denoted as COCO-Binary and consider it as the source dataset, \mathcal{D}_s . More specifically, we generate COCO-Binary by removing all the three categories assigned to the label one, and the rest categories to zero. We then train the objectness network, θ_c , on the source dataset, \mathcal{D}_s , under two different settings which compare a pixel-wise 'objectness score' (similar to the saliency detection model). In the first setting, we include all the images from the source data, \mathcal{D}_s , regardless of whether the objects found in \mathcal{D}_s images overlap with target data, \mathcal{D}_c . In the second setting, we create a subset of \mathcal{D}_s by excluding those images containing any categories which overlap with \mathcal{D}_c categories. We can formalize the overlapping and non-overlapping settings as follows:

$$O = \begin{cases} \mathcal{D}_s & \text{overlapping} \\ \mathcal{D}_s \setminus \mathcal{D}_c & \text{non-overlapping} \end{cases}$$

where \mathcal{D}_c denotes the set of object classes contained in COCO-Binary used to train the objectness model, θ_c . \mathcal{D}_s represents the non-overlapping subset where there is no semantic category overlap between \mathcal{D}_s and \mathcal{D}_c . Note that the semantic annotations are only used to generate the subset of non-overlapping data, $\mathcal{D}_s \setminus \mathcal{D}_c$, and is not required for training the objectness model, θ_c . We believe the non-overlapping setting is more challenging than image-level based WSS [55, 76, 79], because these methods contain semantic overlap within the source and target data. In both settings, we train the objectness classifier using the cross-entropy approximation and use the binary cross-entropy loss function. The main goal of the objectness classifier is to learn a strong objectness prior for the 'thing' label.

Class-Driven Pseudo-Labels

CAM [83] is widely used as a weak source of supervision as it roughly localizes semantic object areas. Following previous works [13], we first generate CAMs for training images by adapting the method of [83] using a multi-label image classification network. For a fair comparison, we use a ResNe50 [24] model as the classification network, we used in other CAM-based methods [2, 3, 76]. We directly utilize these CAMs to generate pseudo-labels by thresholding their confidence scores for each class label at every pixel predicted to be an object by the class-agnostic objectness network (see Fig. 1(b)). We can formalize this procedure as follows:

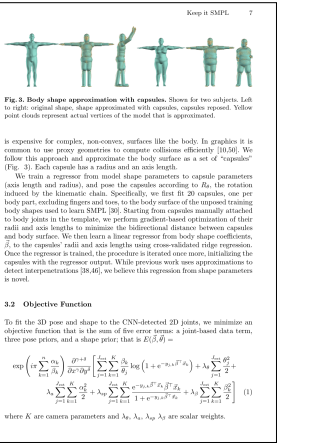
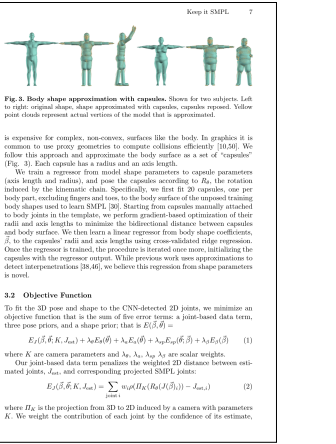
$$p_c = \begin{cases} \max_{i \in \{1, \dots, C\}} \text{CAM}_i(x, y) & \text{if } \text{CAM}_c(x, y) > \tau \\ 0 & \text{otherwise} \end{cases}$$


Figure 1. We applied zero2hero to notoriously unsuccessful papers from the last decade and observe overwhelming results. From left to right: ResiNet [6] paper, a paper on adaptive and robust loss functions [1], *Simpler Does It* [7], and some random work by authors who love to keep it SMPL [2]. Original papers are shown in the top row; note the obscenely simple equations. No wonder these papers have disappeared into oblivion. Professionally over-complicated formulas produced with zero2hero are shown in the bottom row. Smell the success.

the title) or adding (random) symbols, unnecessary operators or made-up arithmetic symbols, any kind of functions (algebraic, arithmetic, barbaric, trigonometric, etc.; see [8] for further information), and/or physical constants (just to name a few). Although these tricks are considered well-known, quite a significant number of submissions are still getting rejected from top-tier conferences every year (which, by the way, causes pain and sorrow across the world). Why the hell is that? How can this even be possible? And, why are we here? With this work, we finally provide complex answers to the big questions¹: Authors simply must have had (and still have) serious problems complicating their papers.

To finally put an end to this tragedy, we propose a fundamentally new method inspired by [10], dubbed zero2hero, that turns every paper into a scientific masterpiece. Based on the latest, next-generation machine-learning techniques, our algorithm methodically over-complicates mathematical equations in a fully automatized way. Just throw in your $\mathcal{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ document, let the machine do the work for you, and et voila, see your paper being accepted at CVPR. No more pain. No more tears. No rejection, no cry! We demonstrate our ground-breaking system, zero2hero, on various notoriously unsuccessful research papers from the last decade, for example, Kaiming He’s ResiNet [6] paper, shown in Fig. 1.

¹We also refer to <https://www.amazon.de/-/en/Stephen-Hawking/dp/1473695988>. Use code ‘SIGBOVIK’ to get 100% off.

2. Methods

Our method assumes as input an ordinary $\mathcal{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ document and outputs a second version of that document where simple equations are replaced by overly-complicated (looking) formulas. It is important to note that that we do *not* care about whether formulas are actually complex, we just want them to look extremely difficult. This is because Attention Is All We Want. Clever, huh?

Following a recent trend kicked off by OpenAI, we do not describe our method due to the competitive nature of our ideas. In particular, we intentionally hide any details about the size of our model, we just want them to look extremely difficult. This is because Attention Is All We Want. Clever, huh?

Following a recent trend kicked off by OpenAI, we do not describe our method due to the competitive nature of our ideas. In particular, we intentionally hide any details about the size of our model, we just want them to look extremely difficult. This is because Attention Is All We Want. Clever, huh?

$$\mathcal{L} = \sum_{i=1}^n \left[-y_i \int_{\Omega} \zeta \left(\frac{\hat{y}_i}{1 - \hat{y}_i} \right) \frac{\partial}{\partial \theta_i} \left(f_i(\theta) \log \frac{\hat{y}_i}{1 - \hat{y}_i} \right) d\theta \right] + \frac{1}{2} \sum_{k=1}^n \frac{\partial^2}{\partial x_k^2} \left(\sum_{i=1}^n y_i \hat{y}_i \frac{\partial \log f_i(\theta)}{\partial x_k} \right), \quad (1)$$

where n is the number of training examples, \hat{h} is the reduced Planck constant, i.e., $\hat{h} = h/2\pi$ with h being the Planck constant, and f_i

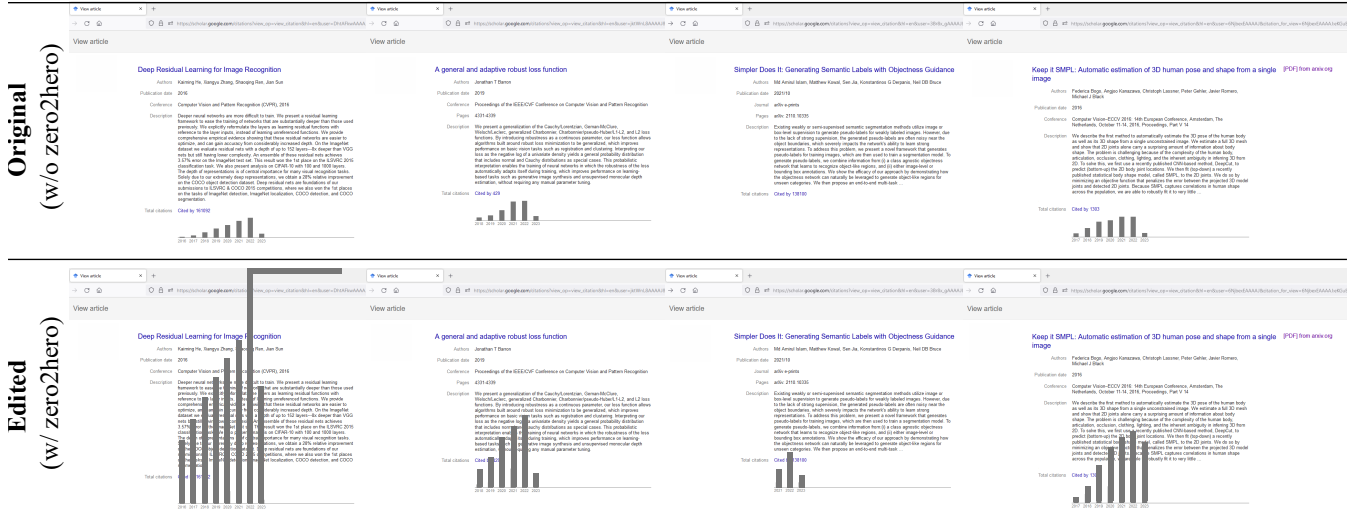


Figure 2. Screenshots of random Google Scholar profiles showing citations of the four investigated papers without using zero2hero (top row) and when zero2hero would have been applied (bottom row) prior to publication. Results are speaking for themselves. Notice how citations of the ResiNet [6] paper (first column) went through the browser bar in 2022. From left to right (same order as in Fig. 1): ResiNet [6] paper, a paper on adaptive and robust loss functions [1], *Simpler Does It* [7], and *Keep It SMPL* [2].

is a secret function transforming parameters θ of the LLM. Please understand that we are not allowed to share any additional details².

3. Experiments and Results

Expensive experiments were conducted to validate our method. Specifically, we analyze two different factors: The impact of zero2hero on (i) the number of citations, and (ii) the author’s mood and personal situation. All experiments were executed retrospectively. Results were analyzed using openCHEAT [4].

3.1 Setup

To analyze how zero2hero would have influenced factors (i) and (ii) for manuscripts written *before* our method was invented, we randomly collected a bunch of papers from the internet and compare the current impact (as of 2023) to what the paper could have had if zero2hero had been used at the time of publication. But, wait, how can we know the impact a paper could have had?

Turns out to be dead easy! In short, to obtain the impact a paper could have generated if zero2hero had been used, we make use of our institution’s high-performance time machine (HPTM) and a theory commonly known as the *many-worlds interpretation*³ (MWI). The MWI is an absurd interpretation of an absurd physics theory (namely, quantum mechanics), asserting that the universal wavefunction is objectively real and that wave functions can’t collapse. This obviously implies that every possible outcome of a decision opens up a new, *parallel universe* (or, world). Given these tools and a paper we want to analyze in our current universe, \mathcal{U}_C , at a certain point in time, t , we proceed as follows.

1. Using our HPTM, we travel back in \mathcal{U}_C to the time shortly before the paper was published (we attached zero2hero to the journey). Denote this point in time as t_0 .

² However, our source code was leaked and submitted to GitHub by a ghost author that we later removed from the planet and the manuscript (in this order).

³ https://en.wikipedia.org/wiki/Many-worlds_interpretation

2. At t_0 , we decide to *not* use zero2hero. Note that (due to the MWI) this decision immediately opens a new universe, \mathcal{U}_N , in which zero2hero is *automatically* applied.
3. In both universes, \mathcal{U}_C and \mathcal{U}_N , we simultaneously publish the paper at time $t_0 + \epsilon$.
4. Lastly, we travel back to where we came from. That was t .

It’s important to note that (in the third step) we do not have to switch the universe in order to publish the paper (in practice, we simply open a new terminal and ssh to \mathcal{U}_N). As such, we never left our current universe.

We now analyze factors (i) and (ii) in detail.

3.2 Impact on Number of Citations

We start by analyzing how zero2hero would have influenced the number of citations for papers written before our method was invented. To do so, the Internet Explorer (version 8.0.7601.17514IC) was used to access Google Scholar profiles from \mathcal{U}_N at time t , again via an ssh connection. We analyze the same four papers shown in Fig. 1, i.e., [6], [1], [7], and [2]. All papers were written and published between 2016 and 2021.

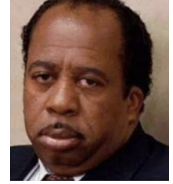
Some exemplary results can be found in Fig. 2. They are clearly out of this world. In all cases, an application of zero2hero would have increased the number of citations dramatically. Most notably, if the authors from *Simpler Does It: Generating Semantic Labels with Objectness Guidance* [7] would have used zero2hero prior to publication in 2021, they could already have 138,100 citations today! Instead, they have zero citations. Well, seems like simpler doesn’t always do it.

3.3 Impact on Mood and Personal Situation

Next, we investigate how zero2hero could have influenced an author’s mood and personal situation. Specifically, we interviewed random people close to an author (family, friends, colleagues) and asked uncomfortable questions about the author’s current personality. As usual, we did this in the current universe \mathcal{U}_C , where the author didn’t use zero2hero as well as in the parallel universe \mathcal{U}_N , where the author did use zero2hero.

w/o zero2hero

"I'm a colleague of ■■■. I got to know ■■■ when I joined his group in 2015. Mr. ■■■ is very talented and clearly loves doing research; unfortunately, as far as I can tell, his career is marked by rejection. In the last three years, almost 65 percent of his papers have been rejected from CVPR. That has not passed him by without leaving a trace. He changed. He looks sad. I wish there's something that could get him back on track... ."



w/ zero2hero

"I've the pleasure to work with ■■■ for 5 years now. Mr. ■■■ is the best boss I've ever had! He's a machine, his papers rock CVPR every single year. I am not absolute certain, yet I believe his unbelievable success stems from his brilliant ability to write and convince with insanely complex papers (I don't know how he's developing all those formulas, he always locks himself when writing papers). I don't understand his manuscripts at all, even though I usually develop the methods about he's writing. That's so cool. He is truly a hero."



Figure 3. Representative result uncovering how zero2hero affects an author's mood and personal situation. The same colleague talking about the same author, however, one time the author didn't make use of zero2hero (top row), and one time he did (bottom row). To respect the author's privacy, we blanked out his name and only show photos (on the right). We clearly see that zero2hero delivers what it promises.

Please find a representing answer from a colleague for one author in Fig. 3. Obviously, as seen, zero2hero has the complex ability to transform people's lives. Sheesh.

4. Limitations

In case our method is applied to an *actually complex equation* (which, luckily, are rather rare and anyway unnecessary in practice) this might overload human brain capacity. Also, do not apply zero2hero multiple times to the same simple equation. **Please consult your doctor or pharmacist if you've overdone it once again** (watch for symptoms such as disorientation or general confusion, in Germany also known as *Verwirtheit* [11]).

Moreover, we do want to note that our implementation of zero2hero may not properly handle complex edge cases and, therefore, might be prone to errors. Due to the severe complexity of zero2hero, however, we do not expect this to be a major limitation in practice as nobody is able to spot those errors anyway. If you do find an issue, please HonkFast [3] and we'll make it work again.

5. Conclusion

It's complicated.

Acknowledgments

Last but not least, I wanna thank me. I wanna thank me for believing in me. I wanna thank me for doing all this hard work. Everyone else: Thanks for nothing.

References

- [1] J. T. Barron. A general and adaptive robust loss function. *CVPR*, pages 4331–4339, 2019.
- [2] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it simple: Automatic estimation of 3d human pose and shape from a single image. *ECCV*, pages 561–578, 2016.
- [3] B. Egger and M. Siegel. Honkfast, prehonk, honkback, prehonkback, hers, adhonk and ahc: the missing keys for autonomous driving. *SIGBOVIK*, 2020.
- [4] B. Egger, K. Smith, and M. Siegel. opencheat: Computationally helped error bar approximation tool-kickstarting science 4.0. *SIGBOVIK*, 2021.
- [5] B. Egger, K. Smith*, T. O'Connell*, and M. Siegel. Action: A catchy title is all you need! *SIGBOVIK*, 2022.
- [6] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CVPR*, pages 770–778, 2016.
- [7] M. A. Islam, M. Kowal, S. Jia, K. G. Derpanis, and N. D. Bruce. Simpler does it: Generating semantic labels with objectness guidance. *BMVC*, 2021.
- [8] K. Smith and B. Egger. (un)helper functions. *SIGBOVIK*, 2022.
- [9] M. Weiherer and B. Egger. A free computer vision lesson for car manufacturers or it is time to retire the erlkönig. *SIGBOVIK*, 2022.
- [10] M. Weiherer and B. Egger. From zero to hero: Convincing with extremely complicated math. *SIGBOVIK (under careful review by very talented, outstanding reviewers)*, 2023.
- [11] V. Wirth. Author-unification: Name-, institution-, and career-sharing co-authors. *SIGBOVIK (under careful review by very talented, outstanding reviewers)*, 2023.