



Theory-based residual neural networks: A synergy of discrete choice models and deep neural networks



Shenhao Wang^{a,*}, Baichuan Mo^a, Jinhua Zhao^a

Massachusetts Institute of Technology, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:

Received 2 April 2020

Revised 12 February 2021

Accepted 3 March 2021

Available online 18 March 2021

Keywords:

Deep learning

Choice modeling

Travel behavior

ABSTRACT

Researchers often treat data-driven and theory-driven models as two disparate or even conflicting methods in travel behavior analysis. However, the two methods are highly complementary because data-driven methods are more predictive but less interpretable and robust, while theory-driven methods are more interpretable and robust but less predictive. Using their complementary nature, this study designs a theory-based residual neural network (TB-ResNet) framework, which synergizes discrete choice models (DCMs) and deep neural networks (DNNs) based on their shared utility interpretation. The TB-ResNet framework is simple, as it uses a $(\delta, 1-\delta)$ weighting to take advantage of DCMs' simplicity and DNNs' richness, and to prevent underfitting from the DCMs and overfitting from the DNNs. This framework is also flexible: three instances of TB-ResNets are designed based on multinomial logit model (MNL-ResNets), prospect theory (PT-ResNets), and hyperbolic discounting (HD-ResNets), which are tested on three data sets. Compared to pure DCMs, the TB-ResNets provide greater prediction accuracy and reveal a richer set of behavioral mechanisms owing to the utility function augmented by the DNN component in the TB-ResNets. Compared to pure DNNs, the TB-ResNets can modestly improve prediction and significantly improve interpretation and robustness, because the DCM component in the TB-ResNets stabilizes the utility functions and input gradients. Overall, this study demonstrates that it is both feasible and desirable to synergize DCMs and DNNs by combining their utility specifications under a TB-ResNet framework. Although some limitations remain, this TB-ResNet framework is an important first step to create mutual benefits between DCMs and DNNs for travel behavior modeling, with joint improvement in prediction, interpretation, and robustness.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

As machine learning (ML) is increasingly used in the transportation field, we observe a tension between data-driven ML methods and classical theory-driven methods. Take travel behavior research as an example: researchers can analyze travel mode choice by using discrete choice models (DCMs) under the framework of random utility maximization (RUM), or using data-driven methods such as ML classifiers without any substantial behavioral understanding. This tension creates practical difficulty in choosing one method over the other, and prevents scholars from tackling travel behavior problems under a unified framework. This tension even delineates a form of partisan line within the transportation research community:

* Corresponding author.

E-mail address: shenhao@mit.edu (S. Wang).

Table 1
Comparison of DNNs, DCMs, and TB-ResNets.

Models	Prediction	Interpretability	Robustness
Deep neural networks (DNNs)	High	Low	Low
Discrete choice models (DCMs)	Low	High	High
Theory-based residual neural networks (TB-ResNets)	High	High	High

researchers using data-driven methods focus on computational perspectives and prediction accuracy, while researchers using theory-driven methods focus on interpretation, economic information, and behavioral foundations.

However, a closer examination reveals that the two methods are complementary in terms of prediction, interpretation, and robustness, prompting us to ask how to synergize them rather than treating them as disparate or even conflicting methods. As summarized in Table 1, deep neural networks (DNNs) and DCMs can be complementary because the former are more predictive, but less interpretable and robust, while the latter are less predictive, but more interpretable and robust. While DNNs are widely known as highly predictive (Kotsiantis et al., 2007; Karlaftis and Vlahogianni, 2011; Krizhevsky et al., 2012; LeCun et al., 2015; Glaeser et al., 2018), researchers often contend that DNNs lack interpretability (Kotsiantis et al., 2007; Lipton, 2016; Doshi-Velez and Kim, 2017), which is crucial in analyzing individual behavior for reasons such as safety in autonomous vehicles, knowledge distillation in research, and transparency in governance (Lipton, 2016; Doshi-Velez and Kim, 2017; Freitas, 2014). DNNs are also found to lack robustness, creating a brittle system vulnerable to small random noises or adversarial attacks. On the other hand, parsimonious DCMs are believed to be more interpretable and robust, although their predictive power can be low owing to their misspecification errors. Therefore, it appears to be a natural question whether these two complementary methods can be synergized to retain the strength from both sides. However, since DNNs and DCMs emerged from two different research communities (computer science and economics), it is unclear whether this synergy is even possible, let alone mutually beneficial.

To address the aforementioned challenge, this study designs a theory-based residual neural network (TB-ResNet) that synergizes DNNs and DCMs, demonstrating that this synergy is not only feasible but also desirable, leading to a simultaneous improvement in prediction, interpretation, and robustness. This study consists of three main components. We first demonstrate that DNNs align with the RUM framework by briefly recounting McFadden (1974) and Wang et al. (2020a). Second, we present the TB-ResNet framework, which augments DNNs to DCMs to fit the utility residuals with a $(\delta, 1 - \delta)$ formulation, resembling the essence of the standard residual network (ResNet) (He et al., 2016). This TB-ResNet framework is further elaborated with six interwoven perspectives: architecture design, model ensemble, gradient boosting, regularization, flexible function approximation, and theory diagnosis. The regularization perspective is formally demonstrated by using the state-of-the-art statistical learning theory to illustrate the intuition that DNNs tend to be too complex to capture the reality and DCMs tend to be too simple to do so. Then we design three instances of TB-ResNets using multinomial logit models (MNL-ResNets), prospect theory (PT-ResNets) for risk preference, and hyperbolic discounting (HD-ResNets) for time preference, showing that the simple TB-ResNet framework can incorporate a wide range of DCMs that are part of the utility maximization framework. Lastly, we use empirical testing to determine whether the three instances of TB-ResNets are effective in three datasets, one collected in Singapore and two from Tanaka et al. (2010). We found that (with some exceptions) the three instances of TB-ResNet can generally improve the overall prediction, interpretability, and robustness of the pure DCMs and DNNs.

The next section reviews related studies. Section 3 introduces the TR-ResNet and its three instances. Section 4 discusses the design of experiments. Section 5 presents the results, and Section 6 concludes and discusses our findings.

2. Literature Review

Individual decision-making has been a classical research question in economics, transportation, marketing, and many other social science and engineering fields. At least three predominant types of DCMs exist: the multinomial logit (MNL) model describing the trade-off between multiple alternatives, prospect theory (PT) models that analyze decision-making under risk and uncertainty, and hyperbolic discounting models that analyze temporal decision-making. McFadden (1974) developed the seminal MNL model based on random utility maximization and applied the model to travel behavior analysis (McFadden, 1974). After McFadden (1974), several generations of researchers refined the MNL model by incorporating heterogeneity, endogeneity, and more complicated substitution patterns (Train, 1980, 2009; Ben-Akiva and Lerman, 1985). In terms of risk preference, Neumann and Morgenstern (1944) created the expected utility model to analyze how individuals make decisions with risky inputs. Kahneman and Tversky (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) created prospect theory (PT), which addresses abnormalities that cannot be explained by the initial expected utility models (Neumann and Morgenstern, 1944; Pratt, 1964; Arrow, 1965; Sydnor, 2010). In the last two decades, researchers gradually improved these models by specifying the formulation of reference points or adding more interactions between attributes and probabilities (Tanaka et al., 2010; Dhami, 2016; Koszegi and Rabin, 2006). With regards to time preference, important models include exponential discounting (Samuelson, 1937), hyperbolic discounting (HD) (Loewenstein and Prelec, 1992), quasi-hyperbolic discounting (O'Donoghue and Rabin, 1999) and many others. Given the ubiquity of individual decision-making across a massive number of fields, these three types of theories have been widely applied to analyze travel behav-

ior, technology adoption, fuel economy, policy decisions, insurance premiums, procrastination, and self-control (Camerer and Kunreuther, 1989; Nicholson and Snyder, 2012; Liu, 2013; O'Donoghue and Rabin, 2001; Kaur et al., 2015).

Recently researchers have started to use DNNs to predict travel behavior, demonstrating that DNNs can outperform DCMs in terms of prediction accuracy, although these studies often fail to connect DNNs to DCMs in a deeper manner. Individual decision-making can be treated as a ML classification task because the target variables are often discrete. Researchers used DNNs to predict travel mode choice (Cantarella and de Luca, 2005), car ownership (Paredes et al., 2017), travel accidents (Zhang et al., 2018b), travelers' decision rules (van Cranenburgh and Alwosheel, 2019), driving behaviors (Huang et al., 2018), trip distribution (Mozolin et al., 2000), hierarchical demand structure (Wu et al., 2018a), queue lengths (Lee et al., 2019), parking occupancy (Yang et al., 2019), metro passenger flows (Hao et al., 2019), and traffic flows (Polson and Sokolov, 2017; Liu and Chen, 2017; Wu et al., 2018; Zhang et al., 2018a; Do et al., 2019; Ma et al., 2020). DNNs are also used to complement the smartphone-based survey (Xiao et al., 2016), improve survey efficiency (Seo et al., 2017), synthesize new populations (Borysov et al., 2019), and impute survey data (Duan et al., 2016). Studies typically found that the ML classifiers, including DNNs, support vector machines, decision trees, and random forests, can achieve higher predictive performance than the classical DCMs (Pulugurta et al., 2013; Omrani, 2015; Sekhar and Madhu, 2016; Hagenauer and Helbich, 2017; Cantarella and de Luca, 2005). However, these investigations are mainly limited to a comparative perspective, implicitly intensifying the tension between the data-driven and the theory-driven methods. While several recent studies started to explore the interaction between DNNs and DCMs (Wang et al., 2018, 2020a, 2020c), the exploration is still inadequate. Given the predominant use of DCMs and DNNs in travel modeling, it is imperative to demonstrate how to adopt the DNN perspective to analyze individual decision-making beyond prediction.

Prediction accuracy should not be the only focus, because interpretability and robustness are both important criteria (Lipton, 2016; Freitas, 2014; Doshi-Velez and Kim, 2017). Although many recent studies have focused on DNN interpretation (Hinton et al., 2015; Ribeiro et al., 2016; Erhan et al., 2009; Baehrens et al., 2010; Szegedy et al., 2014), DNNs are still largely perceived as lacking interpretability (Kotsiantis et al., 2007). This is not surprising given that DNNs were initially designed to maximize the prediction power. DNNs and DCMs respectively focus on prediction and interpretation, or equivalently, on predicting \hat{y} and estimating β , as argued by Mullainathan and Spiess (2017). In the transportation field, only a small number of studies have touched upon the interpretability issue of DNNs for choice modeling. For example, researchers extracted full economic information from DNNs (Wang et al., 2020b), ranked the importance of DNN input variables (Hagenauer and Helbich, 2017), or visualized the input-output relationship to improve the understanding of DNN models (Bentz and Merunka, 2000). The challenge of analyzing model interpretability is partially caused by the ambiguity of its definition, as pointed out by Lipton (2016). For example, interpretability can be defined as “simulatability”: whether researchers can easily simulate the model in their mind. It can also be defined as the capacity to approximate the true probabilistic behavioral mechanism for the choice modeling context (Wang et al., 2020b, 2018). This work adopts the latter definition recognizing the importance of behavioral realism in demand modeling.

In the choice modeling context, robustness represents the local stability of economic information and the regularity of the behavioral mechanism, which is formally measured by the prediction invariance to random noises or adversarial attacks. When a choice model is robust, a small perturbation of the inputs, such as a \$0.1 decrease in transit fare, does not lead to a dramatic change in outputs, such as a dramatic increase in the choice probability of using public transit. While robustness is not a common topic in the classical demand modeling framework, it is becoming more important for the heavily parameterized DNNs, which are more likely to present irregular local patterns and have been widely criticized as lacking robustness (Szegedy et al., 2014; Goodfellow et al., 2015; Wang et al., 2020b). To formally measure robustness, researchers need to examine whether predictive performance significantly decreases under random noises or adversarial attacks. Many adversarial attacks have been created, including the fast gradient sign method (FGSM), one-step target gradient sign method (TGSM), and iterative least-likely class method (Szegedy et al., 2014; Goodfellow et al., 2015; Papernot et al., 2016a,b; Kurakin et al., 2016, 2017). This study will evaluate the robustness of DCMs, DNNs, and TB-ResNets under both the random noises and the adversarial attacks.

3. Theory

Section 3.1 demonstrates that DNNs have an implicit utility interpretation by recounting the results from McFadden (1974) and Wang et al. (2020a). Subsection 3.2 presents the TB-ResNet framework, and introduces six pertinent ML and behavioral perspectives. Subsection 3.3 formally illustrates the regularization perspective to illustrate the rationale underlying the design of TB-ResNets. Subsection 3.4 substantiates this TB-ResNet framework by creating three instances for three choice scenarios.

3.1. Deep neural networks and random utility maximization

The choice analysis includes two types of inputs: alternative-specific variables x_{ik} and individual-specific variables z_i with $i \in \{1, 2, \dots, N\}$ representing the individual index and $k \in \{1, 2, \dots, K\}$ the alternative index. Let $B = \{1, 2, \dots, K\}$ and $\tilde{x}_i = [x_{i1}^T, \dots, x_{iK}^T]^T$. The output is individual i 's choice, denoted as $y_i = [y_{i1}, y_{i2}, \dots, y_{iK}]$, with each $y_{ik} \in \{0, 1\}$ and $\sum_k y_{ik} = 1$. The RUM

framework assumes that individuals maximize the sum of a deterministic utility v_{ik} and a random utility ϵ_{ik} :

$$u_{ik} = v_{ik} + \epsilon_{ik} = V_k(z_i, \tilde{x}_i) + \epsilon_{ik} \tag{1}$$

in which v_{ik} represents the deterministic utility value and V_k the utility function. The probability of individual i choosing alternative k is

$$P_{ik} = \text{Prob}(v_{ik} + \epsilon_{ik} > v_{ij} + \epsilon_{ij}, \forall j \in B, j \neq k) \tag{2}$$

Assuming ϵ_{ik} is independent and identically distributed across individuals and alternatives and the cumulative distribution function of ϵ_{ik} is $F(\epsilon_{ik})$, then

$$P_{ik} = \int \prod_{j \neq k} F_{\epsilon_{ij}}(v_{ik} - v_{ij} + \epsilon_{ik}) dF(\epsilon_{ik}) \tag{3}$$

and the two following propositions hold:

Proposition 1. Suppose ϵ_{ik} follows a Gumbel distribution, with a probability density function equal to $f(\epsilon_{ik}) = e^{-\epsilon_{ik}} e^{-e^{-\epsilon_{ik}}}$ and a cumulative distribution function equal to $F(\epsilon_{ik}) = e^{-e^{-\epsilon_{ik}}}$. Then the choice probability P_{ik} takes the form of the Softmax activation function.

$$P_{ik} = \frac{e^{v_{ik}}}{\sum_j e^{v_{ij}}} \tag{4}$$

Proposition 2. Suppose Equation 3 holds, and choice probability P_{ik} takes the form of the Softmax function as in Equation 4. If ϵ_{ik} is a distribution with the transition complete property, then ϵ_{ik} follows a Gumbel distribution, with $F(\epsilon_{ik}) = e^{-\alpha e^{-\epsilon_{ik}}}$.

Propositions 1 and 2 jointly demonstrate that DNNs have an implicit RUM interpretation. Specifically, the Softmax activation function, which is used in nearly all DNN architectures as the last layer, implies the random utility terms with a Gumbel distribution under the RUM framework. When a fully connected feedforward DNN is applied to inputs \tilde{x}_i and z_i , the implicit assumption is RUM with the random utility term following Gumbel distribution. Therefore, the inputs into the Softmax activation function in DNNs can be interpreted as the utilities of alternatives. The Softmax function itself is a process of comparing utility scores. The DNN transformation prior to the Softmax function is a process of specifying utilities. Proposition 1 can be found in nearly all the textbooks of choice modeling (Train, 2009; Ben-Akiva and Lerman, 1985), and Proposition 2 is from Lemma 2 in McFadden (1974). By taking advantage of the RUM interpretation in DNNs, researchers can design novel DNN architectures to improve the model performance (Wang et al., 2020a). The sketch proof of the two propositions is in Appendix A.

DNNs and DCMs share the utility maximization framework, but they parameterize their utility functions in different ways. DNNs can automatically learn utility function owing to the strong approximation power of complex DNNs' model families (Hornik et al., 1989; Hornik, 1991; Cybenko, 1989), while DCMs rely on much more parsimonious parametric assumptions. For example, the utility function of DNNs ($V_{DNN,k}$) can be parameterized by millions of parameters, while that of DCMs ($V_{T,i}$) often by less than ten parameters. The similar utility interpretation shared by DNNs and DCMs enables us to design the TB-ResNet framework, and their differences in model complexity are an opportunity for complementarity.

3.2. Theory-based residual neural networks

Leveraging the similar utility interpretation in DCMs and DNNs, we design the TB-ResNet framework, which consists of a DCM utility function $V_{T,k}(z_i, \tilde{x}_i)$ and a DNN utility function $V_{DNN,k}(z_i, \tilde{x}_i)$ weighted by δ and $1-\delta$:

$$v_{TB-ResNet,ik} = (1 - \delta)v_{T,ik} + \delta v_{DNN,ik} = (1 - \delta)V_{T,k}(z_i, \tilde{x}_i) + \delta V_{DNN,k}(z_i, \tilde{x}_i) \tag{5}$$

where $V_{T,k}$ represents the utility function from DCMs, $V_{DNN,k}$ represents that from DNNs, and $(\delta, 1-\delta)$ adjusts the weighting between them. This TB-ResNet can be seen as a linear combination of the two types of utility functions with a flexible weighting controlled by δ . The utility specification of a feedforward DNN $V_{DNN,k}$ in Equation 5 can be parameterized as

$$V_{DNN,k}(z_i, \tilde{x}_i) = w_{m,k}^T \Phi(z_i, \tilde{x}_i) = w_{m,k}^T (g_{m-1} \dots \circ g_2 \circ g_1)(z_i, \tilde{x}_i) \tag{6}$$

in which m is the number of layers of DNN, $g_l(t) = \text{ReLU}(W_l^T t + b_l)$, and $\text{ReLU}(t) = \max(0, t)$. The W_l represents the parameters in DNNs for layer l , and the ReLU is one activation function among many alternatives (e.g. Tanh and Sigmoid). The DCM utility specification can be parameterized by various utility theories, which we will discuss in Section 3.4.

Fig. 1 represents the architecture of the TB-ResNet, consisting of the shallow $(1 - \delta)$ DCM and the deep δ DNN blocks for joint specification of the deterministic utility term. Typically the DCM block $(1 - \delta)V_{T,k}$ can be represented by a shallow neural network with one-layer transformation, while the DNN block $\delta V_{DNN,k}$ is represented by a deep structure capable of automatic learning. The DCM and DNN blocks transform the inputs (z_i, \tilde{x}_i) into deterministic utilities, which are further converted into choice probabilities and outputs through the Softmax activation function. This TB-ResNet framework can be understood from six interwoven ML and behavioral perspectives as follows.

First and most intuitively, this TB-ResNet can be seen as a new DNN architecture, because the DCM part in the TB-ResNet represents a shallow neural network and the DNN part represents a deep one. In fact, the name of the TB-ResNet arises

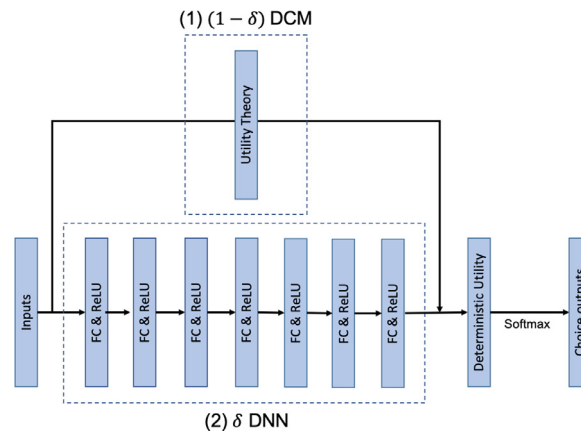


Fig. 1. Architecture of TB-ResNet. Both DCM and DNN are flexible: the DNN block uses seven layers as an example, but it can be any depth or width; the DCM block can take any utility specification under the RUM framework.

from the standard ResNet architecture, which consists of an identity feature mapping and a feedforward DNN architecture: $V_{ResNet,ik} = V_{I,k}(z_i, \tilde{x}_i) + V_{DNN,k}(z_i, \tilde{x}_i)$, where $V_{I,k}(x) = x$. When the true model is close to linear, the ResNet can approximate the true model better than a standard feedforward DNN. This reasoning can be similarly applied to TB-ResNets. When $1 - \delta$ is close to one, the TB-ResNet consists of a main DCM part and a small DNN part fitting the utility residual, resembling the essence of the standard ResNet architecture.

Second, the TB-ResNet framework with the $(\delta, 1 - \delta)$ weighting can be seen as an ensemble model of DCMs and DNNs with scale adjustment. The weighting is controlled by the ratio $\frac{1-\delta}{\delta}$, which can span all possible positive values under a logarithmic scale of $\delta \in (0, 1)$.¹ When $\delta \rightarrow 0$, such as $\delta = 10^{-5}$, the utility ratio $\frac{1-\delta}{\delta}$ converges to $+\infty$; when $\delta \rightarrow 1$, such as $\delta = 1 - 10^{-5}$, the utility ratio $\frac{1-\delta}{\delta}$ converges to 0; when $\delta = 0.5$, the ratio equals to one. Therefore, this $(\delta, 1 - \delta)$ weighting allows us to explore all the possible utility ratio of the DCM and the DNNs. In fact, the flexible scaling is closely related to the randomness discussion in classical choice modeling. For example, to combine revealed and stated preference data, researchers need to adjust a scale factor to reflect the different randomness in two types of data sets (Wang et al., 2020c; Hensher and Bradley, 1993; Bradley and Daly, 1997).

Third, the TB-ResNet framework is similar to the gradient boosting method (Friedman, 2001), although differences still exist. They are similar since both seek to achieve higher performance by adding multiple models; in particular, the TB-ResNets with a sequential training procedure are similar to the boosting method with multiple stages. However, they are also different since gradient boosting is typically used to combine multiple weak classifiers, while TB-ResNets combine a relatively weak classifier (DCMs) and a strong one (DNNs). As a result, the regularization perspective becomes essential in the TB-ResNets, particularly when δ is small. In addition, the TB-ResNets combine the DCMs and DNNs through the shared utility interpretation, while the boosting method connects multiple classifiers through the multi-stage optimization of the loss function. The shared utility interpretation in TB-ResNets is helpful in obtaining not only lower prediction losses, but also improvements in local regularity, robustness, and utility-based economic interpretation. Nonetheless, there exists only a fine line between model ensemble, gradient boosting, and our TB-ResNet framework. It is also possible to improve the TB-ResNet framework by incorporating the other two perspectives in the future.

Fourth, when $\delta \rightarrow 0$, the TB-ResNet is dominated by the DCM component, which becomes a skeleton utility function to localize and stabilize the TB-ResNet system, and the small δ regularizes the complex DNN component to address overfitting. With a small δ , the model complexity between $(1 - \delta)V_{T,k}$ and $\delta V_{DNN,k}$ is more balanced, since the DNN component is typically much more complex than the DCM component. Intuitively, when δ becomes smaller, the TB-ResNet framework is increasingly localized around the DCM component, and the training of the DNN is similar to a searching around the small neighbor of the DCM. A small δ is most effective when simple DCMs can successfully capture much of the true behavioral mechanism. In fact, when the DCM component can perfectly capture the true behavioral mechanism, the best δ should be close to zero.

Fifth, when $\delta \rightarrow 1$, the TB-ResNet is dominated by the DNN component, which allows the TB-ResNet system to use the outstanding approximation power of DNNs to approximate the true data generating process, transcending the theoretical limitations of the DCMs. When δ is closer to one, it trends towards a large regional search by the δ DNN component around the small $(1 - \delta)$ DCM component. A large δ is the most effective when simple DCMs capture little information about the decision-making mechanism. In the worst scenario, when the DCMs capture zero information, the TB-ResNet reduces to a DNN model.

¹ The logarithmic scale refers to the δ taking the values of 10^{-x} and $1 - 10^{-x}$ as it gets close to zero or one. This design is intended to maximize the span of the magnitude of the scale ratio $\frac{1-\delta}{\delta}$.

Therefore, the optimum δ value becomes a metric to diagnose the completeness of the DCMs. A small and optimum δ suggests that the current DCM is highly effective since only a small portion of the DNN component is needed to fit the utility residuals. A large and optimum δ suggests that the current DCM is far from complete since the TB-ResNet mainly uses the DNN component to fit the true behavior. Therefore, the optimum δ value can be a tool to diagnose the completeness of the DCMs. However, the optimum δ value can only be identified empirically since modelers cannot a-priori evaluate the completeness of a theory. The results section will compare the optimum δ values of our three instances, which sheds light on the theoretical completeness of the MNL, PT, and HD models.

In summary, simple DCMs tend to underfit the true behavioral mechanism, while rich DNNs tend to overfit. The TB-ResNets are formulated with a flexible $(\delta, 1 - \delta)$ weighting, taking advantage of both the simplicity of DCMs and the richness of DNNs, and guarding against problems from both sides. A large δ enables the DNN component to fit the utility residual of the DCM component to address the underfitting problem, and a small δ controls the scale of the DNN component as a regularization tool to address the overfitting problem. This is the key intuition underlying the design of TB-ResNets.

3.3. Balancing approximation and estimation errors of DCMs and DNNs

Formally, the problems of underfitting and overfitting can be framed as the challenge of balancing approximation and estimation errors. This subsection will use the state-of-the-art statistical learning theory to illustrate the importance of the δ term in balancing the model complexities between DCMs and DNNs.

The out-of-sample performance of TB-ResNets can be decomposed into approximation and estimation errors, and the analysis into the latter illustrates the importance of controlling model complexity. Let \mathcal{F}_1 and \mathcal{F}_2 denote the model families of DCMs and DNNs. The empirical risk minimization (ERM) is used for model training:

$$\hat{f} = \operatorname{argmin}_{f \in (1-\delta)\mathcal{F}_1 + \delta\mathcal{F}_2} \frac{1}{N} \sum_{i=1}^N l(y_i, f(x_i)) \tag{7}$$

in which x_i is a vector representing both the alternative-specific inputs x_{ik} and the individual-specific inputs z_i . Let f^* denote the true data generating process. The *Excess error* is defined and decomposed as following:

$$\mathbb{E}_S[L(\hat{f}) - L(f^*)] = \mathbb{E}_S[L(\hat{f}) - L(f_{\mathcal{F}}^*)] + \mathbb{E}_S[L(f_{\mathcal{F}}^*) - L(f^*)] \tag{8}$$

where $L = \mathbb{E}_{x,y}[l(y, f(x))]$ is the expected loss function and S represents the sample $\{x_i, y_i\}_1^N$; $f_{\mathcal{F}}^* = \operatorname{argmin}_{f \in \mathcal{F}} L(f)$, the best function in function class $\mathcal{F} := (1 - \delta)\mathcal{F}_1 + \delta\mathcal{F}_2$ to approximate f^* . The excess error measures the average difference of the out-of-sample performance between the estimated function \hat{f} and the true model f^* . The excess error is decomposed as an *estimation error* $\mathbb{E}_S[L(\hat{f}) - L(f_{\mathcal{F}}^*)]$ and an *approximation error* $\mathbb{E}_S[L(f_{\mathcal{F}}^*) - L(f^*)]$. The approximation error is deterministic and is thus irrelevant to the training procedure, so this study does not address approximation errors in details. A simple model family (e.g. DCMs) will usually approximate the true data generating process worse than a complex model family (e.g. DNNs), which can also be inferred from the universal approximator theorem of the DNN (Hornik et al., 1989). The key upper bound is on the estimation error $\mathbb{E}_S[L(\hat{f}) - L(f_{\mathcal{F}}^*)]$, which is provided by using Rademacher complexity from statistical learning theory.

Definition 1. Empirical Rademacher complexity of function class \mathcal{F} is defined as:

$$\hat{\mathcal{R}}_n(\mathcal{F}|_S) = \mathbb{E}_{\epsilon} \sup_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \epsilon_i f(x_i) \tag{9}$$

ϵ_i is the Rademacher random variable, taking values $\{-1, +1\}$ with equal probabilities.

Proposition 3. The estimation error of an estimator \hat{f} can be bounded by the Rademacher complexity of \mathcal{F} .

$$\mathbb{E}_S[L(\hat{f}) - L(f_{\mathcal{F}}^*)] \leq 2\mathbb{E}_S\hat{\mathcal{R}}_n(\mathcal{F}|_S) \tag{10}$$

Definition 1 and **Proposition 3** jointly provide the intuition that the upper bound of the estimation error of any estimate \hat{f} can be approximated by the complexity of the model family \mathcal{F} . Note that $\hat{\mathcal{R}}_n(\mathcal{F}|_S)$ in **Definition 1** measures how complex the model family \mathcal{F} is, and the averaged Rademacher complexity on the right side of **Proposition 3** is the upper bound on the estimation error. In other words, it is important to limit the complexity of the model family \mathcal{F} to achieve a tight upper bound on the estimation error. The proof of **Proposition 3** can be found in Wang et al. (2018), and the key technique is the symmetrization lemma in Wainwright (2019).

In the case of TB-ResNets, \mathcal{F} is designed to be $(1 - \delta)\mathcal{F}_1 + \delta\mathcal{F}_2$. The following three propositions provide the intuition as to why the δ weighting is important:

Proposition 4. The estimation error of the TB-ResNet estimator \hat{f} can be bounded by the weighted Rademacher complexities of \mathcal{F}_1 and \mathcal{F}_2 .

$$\mathbb{E}_S[L(\hat{f}) - L(f_{\mathcal{F}}^*)] \leq 2\mathbb{E}_S[(1 - \delta)\hat{\mathcal{R}}_n(\mathcal{F}_1|_S) + \delta\hat{\mathcal{R}}_n(\mathcal{F}_2|_S)] \tag{11}$$

Proposition 5. The Rademacher complexity of the DCM model family \mathcal{F}_1 can be bounded by

$$\hat{\mathcal{R}}_n(\mathcal{F}_1|_S) \lesssim O\left(\sqrt{\frac{\nu}{N}}\right) \tag{12}$$

in which ν is the VC dimension of function class \mathcal{F}_1 and N is the sample size.

Proposition 6. The Rademacher complexity of the DNN model family \mathcal{F}_2 can be bounded by

$$\hat{\mathcal{R}}_n(\mathcal{F}_2|_S) \lesssim \frac{(\sqrt{2 \log(D)} + 1) \sqrt{\frac{1}{N} \sum_{i=1}^N \|x_i\|^2}}{\sqrt{N}} \times \prod_{j=1}^D M_F(j) \tag{13}$$

in which D represents the depth of DNN, $\|x_i\|$ represents the norm of input variables, and $M_F(j)$ is the upper bound of the Frobenius norm of parameter matrix W_j . Here the DNN model uses ReLU activation functions.

Propositions 4, 5, and 6 demonstrate the importance of using δ to balance the estimation errors between \mathcal{F}_1 and \mathcal{F}_2 . Since the DNN model family \mathcal{F}_2 is much more complicated than the DCM model family \mathcal{F}_1 , a strong regularization needs to be imposed on the DNN model family \mathcal{F}_2 to guarantee a balanced complexity between $(1 - \delta)\hat{\mathcal{R}}_n(\mathcal{F}_1|_S)$ and $\delta\hat{\mathcal{R}}_n(\mathcal{F}_2|_S)$. Specifically, the DNN complexity, as shown in Proposition 6, is exponential in the depth D and also depends on the size of the input variables x_i . On the other hand, the DCM model \mathcal{F}_1 only depends on the VC dimension, which is often linearly related to the number of parameters. Therefore, the DNN complexity is typically much larger than the DCM complexity ($\hat{\mathcal{R}}_n(\mathcal{F}_2|_S) \gg \hat{\mathcal{R}}_n(\mathcal{F}_1|_S)$). In this case, a small δ weighting on the DNN part can limit the total complexity of the TB-ResNets ($(1 - \delta)\mathcal{F}_1 + \delta\mathcal{F}_2$), thus improving the model performance. The proof of Proposition 5 can be found in Wang et al. (2020b), and that of Proposition 6 is in Golowich et al. (2017). We provide a sketch proof for Proposition 4 in Appendix D.

3.4. Three instances of TB-ResNets

3.4.1. Multinomial logit based residual neural networks (MNL-ResNets)

The TB-ResNet framework is flexible because the DCM utility $V_{T,k}$ can take different forms depending on the context. In the MNL setting, we use the linear utility specification as the theory part of TB-ResNets. Replacing the $V_{T,k}$ in Equation 5 by $V_{MNL,k}$ and removing index i for simplicity, the MNL-ResNet is formulated as

$$v_{MNL-ResNet,k} = (1 - \delta)v_{MNL,k} + \delta v_{DNN,k} = (1 - \delta)V_{MNL,k}(z, \tilde{x}) + \delta V_{DNN,k}(z, \tilde{x}) \tag{14}$$

$$V_{MNL,k}(z, \tilde{x}) = w_{0,k} + w'_{x_k} x_k + w'_z z \tag{15}$$

where w_{x_k} represents the parameters for the alternative-specific variables and w_z represents the parameters of the individual-specific variables. This linear-in-parameter specification of $V_{MNL,k}$ is widely used in choice modeling.

3.4.2. Prospect theory based residual neural networks (PT-ResNets)

In the risk preference setting, we use prospect theory as the theory part of TB-ResNets. Replacing $V_{T,k}$ in Equation 5 by $V_{PT,k}$, the utility specification of PT-ResNet is formulated as

$$v_{PT-ResNet,k} = (1 - \delta)v_{PT,k} + \delta v_{DNN,k} = (1 - \delta)V_{PT,k}(z, \tilde{x}) + \delta V_{DNN,k}(z, \tilde{x}) \tag{16}$$

$$V_{PT,k}(z, \tilde{x}) = \sum_j c(x_{kj}) \pi(p_{kj}) \tag{17}$$

Then the value function $c(x_{kj})$ and probability weighting function $\pi(p_{kj})$ are further parameterized as

$$c(x_{kj}) = \begin{cases} x_{kj}^r & x_{kj} \geq 0 \\ -\lambda(-x_{kj})^r & x_{kj} \leq 0 \end{cases} \tag{18}$$

$$\pi(p_{kj}) = e^{-(-\ln p_{kj})^\alpha} \tag{19}$$

$$\alpha = \alpha(z) = \alpha_0 + z'w_{\alpha z} \tag{20}$$

$$r = r(z) = r_0 + z'w_{rz} \tag{21}$$

$$\lambda = \lambda(z) = \lambda_0 + z'w_{\lambda z} \tag{22}$$

In the equations above, j is the index of uncertain monetary payoffs; x_{kj} is the monetary payoff for alternative k at the value indexed by j ; p_{kj} is the winning probability of x_{kj} ; α represents the probability weighting factor; r represents the concavity of value functions; λ represents the loss aversion factor; α , r , and λ are individual specific and can be partially explained by socioeconomic variables z . The specification in Equations from 14 to 22 is basically the same as Tanaka et al. (2010).² PT is widely used for travel behavioral analysis because travel decisions often involve time uncertainty (De Palma et al., 2008; Avineri and Bovy, 2008).

3.4.3. Hyperbolic discounting based residual neural networks (HD-ResNets)

Another instance of TB-ResNets is the hyperbolic discounting residual network (HD-ResNet) for time preferences. The utility function of the HD-ResNet is formulated as following:

$$v_{HD-ResNet,k} = (1 - \delta)v_{HD,k} + \delta v_{DNN,k} = (1 - \delta)V_{HD,k}(z, \tilde{x}) + \delta V_{DNN,k}(z, \tilde{x}) \quad (23)$$

$$V_{HD,k}(z, \tilde{x}) = \sum_j x_{kj} \beta e^{-rt_{kj}} \quad (24)$$

Then β and r can be further parameterized as

$$\beta = \beta(z) = \beta_0 + z'w_{\beta z} \quad (25)$$

$$r = r(z) = r_0 + z'w_{rz} \quad (26)$$

In the equations above, x_{kj} is the monetary payoff; t_{kj} is the associated time; r is the conventional time discounting factor; β is the present-bias factor; both r and β can be partially explained by socioeconomic variables z . The specifications from Equation 23 to 26 are the same as Tanaka et al. (2010).

4. Experiment Setup

4.1. Datasets

The experiments use three datasets for the three instances. The first data set (SG) is a stated preference survey collected by the authors in Singapore in 2017, focusing on the travel mode choice between walking, buses, ridesharing, driving, and autonomous vehicles. The second and third datasets are stated preference surveys collected from Tanaka et al. (2010), which focused on the risk and time preference of two monetary alternatives.

The summary statistics of the three data sets are provided in Appendix B. The sample sizes of the three datasets are respectively 8,418, 6,335 and 5,340, which are similar to typical travel behavioral surveys. The choices in the three data sets are all balanced: even in the most unbalanced case, the smallest share is the 10.38% of the respondents who chose to walk in Singapore. The attributes in the SG data set were designed by using the standard orthogonal experiment design based on the average travel time and costs of the travel alternatives in Singapore. The attributes in the PT and HD data sets were designed by using simulations based on PT and HD theories. The survey for the SG data set was conducted with the help of an online company Qualtrics; the surveys for the PT and HD data sets were collected through interviews with the help of local officials. For data collection details of the three data sets, readers can refer to Wang and Zhao (2019) and Tanaka et al. (2010).

4.2. Training

The three data sets are divided into training and testing sets, with the former used to train the TB-ResNets and the latter to demonstrate the effects of different δ values. The sample sizes of the training and testing sets in the SG, PT, and HD data sets are (5050, 1648), (4199,1050), and (4272, 1068), respectively. To concentrate the focus on the δ factor, we largely simplified the hyper-parameter searching of the DNN component by using only two simple DNN architectures. The first DNN architecture is designed with depth = 3, width = 100, number of iterations = 5,000, and size of mini-batch = 100, and the second DNN architecture is the same as the first except with depth = 5³. The two DNN architectures are compared to reveal that the TB-ResNets performance partially depends on the configuration of the DNN part. However, this work mainly seeks to demonstrate the benefits through the synergistic TB-ResNet framework rather than the improvements *within* either the DCM or the DNN component. It is because either the DCM or the DNN component can be improved by a vast number

² There are two slight differences between our PT model and Tanaka et al. (2010): the initial paper used a non-parametric method to estimate individuals' risk preference parameters and sequentially estimated the coefficients $w_{\alpha z}$, w_{rz} , and $w_{\lambda z}$, while our PT model follows a parametric method and simultaneously estimates all the coefficients.

³ The hyperparameters can vary; for example, the mini-batch size is typically specified as a power of two (e.g. 64, 128, etc.) rather than 100. Our experiments show that the results are robust to different hyperparameter specifications.

of methods, given the past four decades' work on the DCMs and the recent popularity of the DNNs. Our empirical result seeks to demonstrate that, conditioning on any given DCM or DNN, the TB-ResNets can generate mutual benefits for both components.

The effects of δ values are demonstrated in the testing sets and explored on a logarithmic scale.⁴ A small δ such as 10^{-5} postulates a large utility ratio ($\frac{1-\delta}{\delta}$) between the DCM and the DNN parts as 10^5 , while a large δ such as $1 - 10^{-5}$ postulates a small ratio as 10^{-5} . The range of δ values can assist in empirically evaluating the completeness of the DCM theory, which is impossible to know a priori.⁵ An empirically optimum and small δ suggests the relative completeness of the DCM theory, while an optimum and large δ suggests that the DCM model is far from complete and needs further development. The optimality of the δ values will be discussed in all three instances of the TB-ResNets in our result section.

The three instances of TB-ResNets can be trained in a simultaneous or sequential manner. The simultaneous training implies the training of $V_{DNN,k}$ and $V_{T,k}$ at the same time with the stochastic gradient descent. The sequential training implies a sequential approach of training $V_{T,k}$ on the first stage and training $V_{DNN,k}$ on the second. While both methods are consistent with our ML and behavioral perspectives in the TB-ResNet framework, we tend to believe that sequential training is more intuitive than simultaneous training. This is because the DNN component in the TB-ResNet might easily capture all the valuable information that could have been explained by $V_{T,k}$, when $V_{T,k}$ and $V_{DNN,k}$ are simultaneously trained. In this case, the simultaneous training might damage the capacity of $V_{T,k}$ to stabilize the utility function. Nonetheless, we provide the results of both sequential and simultaneous training in our results section and appendices.

By using w_T to represent the parameters in the DCM part and w_{DNN} to represent the parameters in the DNN part, the formulation of the sequential training procedure is described as follows. In the first stage, empirical risk minimization (ERM) is formulated as shown:

$$\min_{V_{T,k} \in \mathcal{F}_1} L(\tilde{x}_i, z_i; w_T) = \min_{V_{T,k} \in \mathcal{F}_1} -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K y_{ik} \log \frac{e^{(1-\delta)V_{T,k}(\tilde{x}_i, z_i; w_T)}}{\sum_{j=1}^K e^{(1-\delta)V_{T,j}(\tilde{x}_i, z_i; w_T)}} \quad (27)$$

which is the same as the maximum likelihood estimation used in classical choice models. The second stage is another training conditioned on \hat{w}_T . With \mathcal{F}_2 representing the model family of DNNs, second-stage training is:

$$\min_{V_{DNN,k} \in \mathcal{F}_2} L(\hat{w}_T, \tilde{x}_i, z_i; w_{DNN}) = \quad (28)$$

$$\min_{V_{DNN,k} \in \mathcal{F}_2} -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K y_{ik} \log \frac{e^{(1-\delta)\hat{V}_{T,k}(\tilde{x}_i, z_i) + \delta V_{DNN,k}(\tilde{x}_i, z_i; w_{DNN})}}{\sum_{j=1}^K e^{(1-\delta)\hat{V}_{T,j}(\tilde{x}_i, z_i) + \delta V_{DNN,j}(\tilde{x}_i, z_i; w_{DNN})}} \quad (29)$$

To evaluate predictive performance, this work uses three metrics: prediction accuracy, cross-entropy loss, and F1 score. Prediction accuracy is formulated as:

$$Accuracy = 1 - \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\hat{y}_i \neq y_i\} \quad (30)$$

in which \hat{y}_i represents the vector of the predicted choice, $\mathbb{1}\{\}$ is an indicator function, and the inequality sign implies that the two vectors (\hat{y}_i and y_i) are different. The cross-entropy loss is formulated as:

$$Cross\ entropy\ loss = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K y_{ik} \log \hat{P}_{ik} \quad (31)$$

in which \hat{P}_{ik} is the predicted choice probability from DCMs, DNNs, or TB-ResNets. This cross-entropy loss is also named as log-loss, the same as the negative value of the log likelihood. Lastly, the F1 score is:

$$F1\ score = \sum_{k=1}^K \mathbb{W}_k \cdot 2 \times \frac{Precision_k \times Recall_k}{Precision_k + Recall_k} \quad (32)$$

where \mathbb{W}_k is the share of label k in the sample ($\mathbb{W}_k = \frac{1}{N} \sum_{i=1}^N y_{ik}$). In the equation above, $Precision_k$ and $Recall_k$ are class-specific and formulated as:

$$Precision_k = \frac{\sum_{i=1}^N \mathbb{1}\{\hat{y}_{ik} = y_{ik}\}}{\sum_{i=1}^N \hat{y}_{ik}} \quad (33)$$

⁴ The δ spans a list of values: [1e-10, 1e-8, 1e-7, 1e-6, 1e-5, 1e-4, 0.001, 0.002, 0.004, 0.005, 0.006, 0.007, 0.008, 0.009, 0.01, 0.03, 0.05, 0.1, 0.3, 0.5, 0.8, 0.9, 0.95, 0.99, 0.999, 0.9999, 1]

⁵ For application, the authors suggest dividing the datasets into training, validation, and testing sets. The validation set can be used to choose the optimum δ value and the testing set for application. Since this work focuses on demonstrating the effects of δ , we decided to simplify the data processing into the division of training and testing sets.

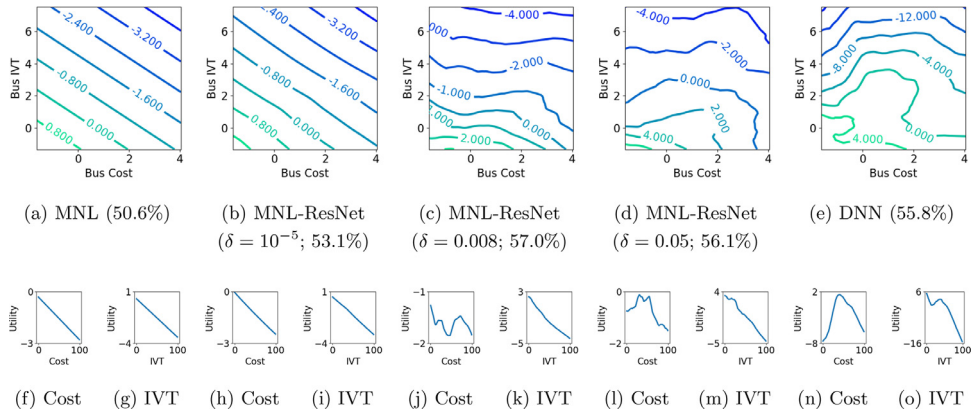


Fig. 2. Utility functions of MNL-ResNets, MNL, and DNNs. Upper row: visualization of 2D utility functions, and percentages in the parentheses represent the prediction accuracy. Lower row: visualization of 1D utility functions, and every pair of figures on the lower row correspond to the figure directly above on the upper row.

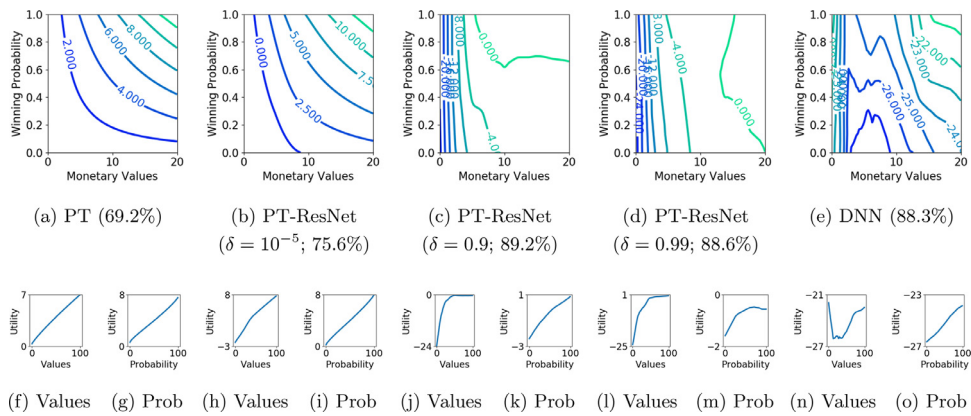


Fig. 3. Utility functions of PT-ResNets, PT, and DNNs (Same format as above)

$$Recall_k = \frac{\sum_{i=1}^N \mathbb{1}\{\hat{y}_{ik} = y_{ik}\}}{\sum_{i=1}^N y_{ik}} \quad (34)$$

Intuitively, $Precision_k$ and $Recall_k$ measure the column- and row-specific performance of the confusion matrix. The F1 score combines the column- and row-specific perspectives with a class-specific weighting. Our three metrics, including accuracy, cross-entropy loss, and F1 score, have taken into account the deterministic and probabilistic decision rules for both balanced and unbalanced outputs.

5. Results

The result section starts with visualization of the utility functions of DCMs, DNNs, and TB-ResNets, thus providing intuition into why DCMs and DNNs are complementary and how TB-ResNets strike a middle ground that retains benefits from both sides. Then we will successively delve into interpretability, prediction, and robustness for comparison and evaluation.

5.1. Utility functions of TB-ResNets as combination of DCMs and DNNs

Figs. 2, 3, and 4 visualize how utilities vary with input values in the three choice scenarios. Take Fig. 2 as an example. The five graphs on the upper row represent how the utility of taking buses varies jointly with two dimensions, the monetary cost (x-axis) and the in-vehicle travel time (y-axis), and the ten graphs on the lower row visualize how utility varies with single dimensions, the monetary cost and the in-vehicle travel time respectively, holding all the other variables constant. Each pair of graphs on the lower row correspond to the one graph directly above them on the upper row. On the upper row of Fig. 2, the figure on the right end visualizes the utility function of DNN; on the left is the utility function of MNL; the three in the middle visualize the utility functions of MNL-ResNets with different δ values. The format of Figs. 3 and 4 is

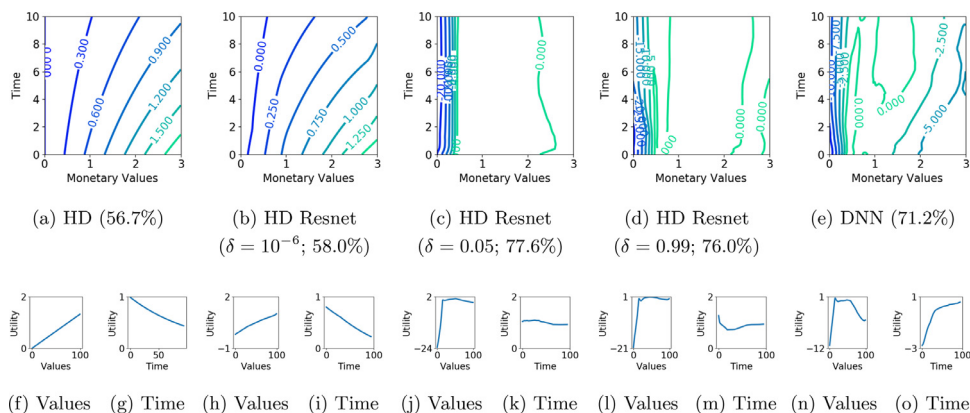


Fig. 4. Utility functions of HD-ResNets, HD, and DNNs (Same format as above)

similar to Fig. 2. In Fig. 3 (PT), the monetary payoff and winning probabilities are the x- and y-axes on the upper row and the x-axes on the lower row. In Fig. 4 (HD), the monetary payoff and time are the x- and y-axes on the upper row and the x-axis on the lower row. All the models in the three figures are trained by using the sequential training method and with the three-layer DNN part. The results of simultaneous training are attached in Appendix C and that of the five-layer DNN are attached in Appendix E. Both yield findings similar to the results in this subsection.

First, we can observe the complementary nature of DNNs and DCMs by comparing only the two graphs of DNNs and DCMs on the right and left ends of Figs. 2, 3, and 4. On one hand, the utility functions of the MNL, PT, and HD models are very regular and intuitive, as shown by subfigures 2a, 3a, and 4a. In subfigures 2f and 2g, the utility values of choosing the bus mode linearly decrease as bus costs and in-vehicle travel time increase. In subfigures 3d and 3e, the utility of taking the risky alternative increases as the monetary payoff and winning probabilities increase. These highly regular utility functions in DCMs are interpretable, although it is also likely that the true utility functions can be much more complex than the smooth and regular MNL, PT, and HD, leading to their misspecification errors and underfitting. However, on the other hand, the utility functions of DNNs for the MNL, PT, and HD scenarios are very irregular and highly counter-intuitive, although they have higher prediction accuracy, as shown by subfigures 2e, 3c, and 4c. For example, in Fig. 2n, the DNN predicts that the utility of using buses first increases as the travel cost increases, violating the basic principle of economics theory. The same type of counter-intuitive results also arise from DNNs in the PT and HD scenarios, as shown in subfigures 3c and 4c. These highly irregular utility functions in DNNs are not very interpretable, although the overly complex functions of DNNs may capture more behavioral mechanisms than DCMs, leading to higher prediction accuracy. In fact, DNNs do outperform DCMs in all three scenarios, as the prediction accuracy of DNN in the MNL setting is 55.8%, higher than 50.6% of MNL; that of DNN in the PT setting is 88.3%, higher than 69.2% of PT; that of DNN in the HD setting is 71.2%, higher than 56.7% of HD. Overall, it is critical to observe the complementary nature of DNNs and DCMs: DCMs might be too simple and regular to capture reality, while DNNs might be too complex and irregular to do so.

TB-ResNets achieve a flexible compromise between DCMs and DNNs, the degree of which is controlled by δ . Take the MNL setting (Fig. 2) as an example. As δ increases, the utility function of MNL-ResNets becomes more irregular and similar to the DNN model, and as δ decreases, the utility function becomes more regular and thus similar to the MNL model. This compromise also happens to the PT and HD settings: PT-ResNets and HD-ResNets resemble a continuum between the highly regular DCMs and irregular DNNs, with δ as the weighting factor.

This perspective of TB-ResNets acting as a flexible compromise between DCMs and DNNs is tied to the six perspectives that were previously introduced. It should be obvious now that this TB-ResNet framework seeks to strike a weighted ensemble model between DCMs and DNNs through the shared utility interpretation, as shown by the continuous change in Figs. 2, 3, and 4. When $\delta \rightarrow 0$, the TB-ResNets increasingly resemble the theory-driven models. The DCM part becomes the skeleton utility function to stabilize the full TB-ResNet utility function, and the DNN part is strongly regularized around the DCM part. When $\delta \rightarrow 1$, the TB-ResNets can gain higher prediction accuracy, and the behavioral patterns become more similar to the irregular DNNs. Only in the middle ground, can the optimum δ values be found to construct the optimum TB-ResNets.

5.2. Interpretability

DCMs tend to be too simple to capture reality, while DNNs tend to be too complex to do so. As a result, TB-ResNets are more interpretable than pure DCMs because TB-ResNets enrich the overly simple utility functions of DCMs with the DNN component, and also more interpretable than the pure DNNs because TB-ResNets regularize the overly complex DNNs with a small δ . For example, the MNL-ResNet model ($\delta = 0.008$ in Fig. 2c) is similar to the MNL model, since it has a relatively regular utility contour and the overall utility values decrease as the cost and in-vehicle travel time increase. But unlike the MNL model, the MNL-ResNet ($\delta = 0.008$) has a richer pattern to reflect the real decision-making mechanism. On the

Table 2
Elasticity coefficients of MNL, DNN, and MNL-ResNet ($\delta = 0.008$).

Panel 1: MNL	Walk	Bus	Ridesharing	Drive	AV
Walk: walk time	-1.778	0.127	0.127	0.127	0.127
Bus: cost	0.149	-0.577	0.149	0.149	0.149
Bus: in-vehicle time	0.121	-0.431	0.121	0.121	0.121
Ridesharing: cost	0.024	0.024	-0.192	0.024	0.024
Ridesharing: in-vehicle time	0.089	0.089	-0.784	0.089	0.089
Drive: cost	0.274	0.274	0.274	-0.760	0.274
Drive: in-vehicle time	0.260	0.260	0.260	-0.418	0.260
AV: cost	0.044	0.044	0.044	0.044	-0.408
AV: in-vehicle time	0.058	0.058	0.058	0.058	-0.540
Panel 2: DNN	Walk	Bus	Ridesharing	Drive	AV
Walk: walk time	-6.298	1.572	1.022	0.292	0.597
Bus: cost	-0.563	-1.842	-0.193	0.739	1.505
Bus: in-vehicle time	0.641	-2.152	-0.494	0.296	-0.184
Ridesharing: cost	0.604	0.514	-2.553	0.062	2.051
Ridesharing: in-vehicle time	-0.534	-0.832	-5.236	0.553	0.601
Drive: cost	0.502	1.557	2.777	-2.133	1.984
Drive: in-vehicle time	1.041	2.787	2.257	-1.643	3.130
AV: cost	-0.828	-0.320	0.841	0.388	-5.919
AV: in-vehicle time	-0.478	0.752	0.794	0.041	-4.594
Panel 3: MNL-ResNet	Walk	Bus	Ridesharing	Drive	AV
Walk: walk time	-2.476	0.478	0.375	0.189	0.275
Bus: cost	-0.262	-0.936	-0.017	0.420	0.346
Bus: in-vehicle time	0.363	-0.785	0.035	0.145	0.488
Ridesharing: cost	-0.108	0.099	-1.003	0.043	0.627
Ridesharing: in-vehicle time	-0.117	-0.109	-1.472	0.210	0.056
Drive: cost	0.283	0.632	0.708	-0.908	0.486
Drive: in-vehicle time	0.465	0.802	0.569	-0.880	0.569
AV: cost	-0.010	0.122	0.481	0.195	-1.670
AV: in-vehicle time	-0.342	0.232	0.289	0.190	-1.783

other side, compared to pure DNNs, the MNL-ResNet ($\delta = 0.008$) retains the general decreasing trend in the utility function, capturing the pattern that the utility of taking buses should decrease with higher travel time and costs. The MNL-ResNet ($\delta = 0.008$) is more interpretable than the DNN owing to the regular local utility pattern in the base MNL component.

PT-ResNets and HD-ResNets demonstrate the same story. An axiom underlying PT is its monotonic increasing property, since larger winning rates and monetary payoffs lead to a higher probability of choosing the alternative. However, this monotonic increasing property is violated in the DNN model, as shown in Figs. 3e, 3n, and 3o, rendering the DNN model difficult to interpret. Compared to pure DNNs, the optimum PT-ResNet ($\delta = 0.9$ in Fig. 3b) retains the monotonicity of the PT model and improves it by allowing a richer pattern augmented by the DNN component. In terms of the HD models, HD-ResNets reveal a similarly successful pattern, showing richer utility patterns than the pure HD and more regular patterns than the pure DNN. The utility function of the HD model is supposed to increase with higher monetary value and decrease with longer waiting time (more temporal discounting), as shown in Figs. 4a, 4d, and 4e. The best HD-ResNet ($\delta = 0.05$ in Figure b) retains these overall patterns with richer details, and is more reasonable than the DNN model (Figs. 4c, f, and 4g), which produced the counter-intuitive result that the utility of a future payoff increases with longer waiting time.

Elasticity is another important metric used to interpret choice models. Although the utility functions of DNNs and TB-ResNets are highly nonlinear, the most practical elasticity values are how the aggregate choice probabilities (market shares) respond to a 1% change in the cost, which can simulate the policy scenarios of imposing a gasoline tax or increasing bus fares. The elasticity can be computed by the formula:

$$\sum_{i=1}^N \frac{\partial P_{i,k_1}/P_{i,k_1}}{\partial x_{i,k_2}/x_{i,k_2}} = \sum_i \frac{\partial P_{i,k_1}}{\partial x_{i,k_2}} \times \frac{x_{i,k_2}}{P_{i,k_1}} \quad (35)$$

in which k_1 and k_2 are the indices of two alternatives and i the index of individuals. When $k_1 = k_2$, this equation computes the self-elasticity; when $k_1 \neq k_2$, it computes the cross-elasticity. With this formula, the elasticity can be computed for every output alternative regarding every input.

Table 2 summarizes the elasticity coefficients of MNL, DNNs, and MNL-ResNet with respect to the alternative-specific variables, demonstrating again that MNL-ResNets achieve a reasonable compromise between MNL and DNN. As shown in Panel 1, the elasticities of MNL follow the independence of irrelevant alternative (IIA) assumption, as the coefficients of self-elasticities are negative and that of cross-elasticities are positive. This pattern is simple and intuitive, but is often criticized as being too restrictive. On the other side, the DNN model reveals a much more irregular elasticity pattern, in which the coefficients of self-elasticities and many cross-elasticities are of a large magnitude and are negative. The elasticity coefficients

Table 3
Performance of DCMs, DNNs, and TB-ResNets in testing sets (Sequential training).

	Prediction accuracy (3-layer)	Cross-entropy loss (3-layer)	F1 score (3-layer)	Prediction accuracy (5-layer)	Cross-entropy Loss (5-layer)	F1 score (5-layer)	Baseline (Largest share)
Panel 1. Performance of MNL models							
MNL	50.6%	1.254	0.439	50.6%	1.254	0.439	44.8%
MNL ResNet ($\delta = 1e-5$)	53.1%	1.207	0.485	52.1%	1.348	0.498	44.8%
MNL ResNet ($\delta = 0.008$)	57.0%	1.237	0.559	56.4%	3.118	0.557	44.8%
MNL ResNet ($\delta = 0.05$)	56.1%	1.572	0.559	57.8%	3.342	0.576	44.8%
DNN	55.8%	2.861	0.555	55.3%	4.175	0.548	44.8%
Panel 2. Performance of PT models							
PT	69.2%	0.602	0.693	69.2%	0.602	0.693	53.8%
PT ResNet ($\delta = 1e-05$)	75.6%	0.502	0.762	81.3%	0.477	0.772	53.8%
PT ResNet ($\delta = 0.9$)	89.2%	0.343	0.887	89.8%	0.366	0.900	53.8%
PT ResNet ($\delta = 0.99$)	88.6%	0.318	0.882	90.0%	0.337	0.899	53.8%
DNN	88.3%	0.353	0.885	89.0%	0.368	0.891	53.8%
Panel 3. Performance of HD models							
HD	56.7%	0.684	0.568	56.7%	0.684	0.568	50.0%
HD ResNet ($\delta = 1e-05$)	68.9%	0.523	0.689	73.7%	0.471	0.737	50.0%
HD ResNet ($\delta = 0.05$)	77.6%	0.437	0.764	75.4%	0.445	0.767	50.0%
HD ResNet ($\delta = 0.99$)	76.0%	0.444	0.763	74.9%	0.453	0.760	50.0%
DNN	71.2%	0.909	0.722	72.6%	0.958	0.726	50.0%

in the MNL-ResNet achieve a certain compromise between MNL and DNN because the magnitude of the coefficients in MNL-ResNet is much smaller than DNNs and shrinks toward the magnitude of the MNL model, and the MNL-ResNet still retains a relatively flexible substitution pattern as in the DNN.

However, the concept of interpretability is ambiguous, since it has multiple definitions, including simulatability, decomposability, algorithmic transparency, and post-hoc interpretability (Lipton, 2016). Our evaluation above has defined interpretability as accurately approximating the true behavioral mechanism and being consistent with the accepted behavior knowledge (Wang et al., 2018, 2020b), which belongs to one type of post-hoc interpretability (Ribeiro et al., 2016; Montavon et al., 2018; Hinton et al., 2015). Mathematically, it is defined as the distance between the true and estimated choice probability functions (Wang et al., 2018), since an accurately estimated choice probability function can provide the complete economic information in demand modeling (Wang et al., 2020b). However, a different definition of interpretability can lead to a different conclusion. For example, suppose that interpretability is defined as simulatability, which means the ease with which modelers can “simulate” the model structure in their mind. The simplest DCMs are then fully interpretable while complex DNNs are not. The TB-ResNet, as a mixture of DCMs and DNNs, can hardly be more interpretable than simple DCMs. Although intuitive, this simulatability definition can be misleading, since it encourages an overly simplified model that fails to recognize rich behavioral reality. For example, a constant model stripped of all the contents (e.g. $y = 0$) can be evaluated as the most interpretable, although it is not practically useful. Interpretability can also be defined by using sensitivity analysis, numerically computing elasticity values for model interpretation (Wang et al., 2020b). Overall, it is important to recognize that many definitions of interpretability can co-exist, from which we adopted only a single relevant one for this work.

5.3. Prediction

Table 3 summarizes the predictive performance of the three choice scenarios in three panels. It presents three metrics - prediction accuracy, cross-entropy loss, and F1 score - to measure the performance of two TB-ResNets with different DNN architectures, one with three layers and the other with five. Each panel includes the DCMs, DNNs, and TB-ResNets with varying δ values. Columns 2-6 present the prediction accuracy, cross-entropy loss, and F1 score in the testing sets of the two architectures, and the last column reports the largest share of the choice set as a baseline performance metric. The convergence of the training process is demonstrated in Appendix F. To provide an intuition for the predictive performance, Fig. 5 visualizes how the model performance varies with δ in the three choice scenarios, with the points representing the individual model performance, fitted by smooth curves, and the red dash lines marking the optimum δ values.




Fig. 5. Model performance and delta (TB-ResNets with the 3-layer DNN)

MNL-ResNets, PT-ResNets, and HD-ResNets outperform both DNNs and DCMs in prediction accuracy, cross-entropy losses, and F1 score, as shown in Table 3 and Fig. 5. In Fig. 5, the curves of the prediction accuracy are always concave and that of the cross-entropy loss are always convex. It suggests that TB-ResNets can outperform both the right and left ends in predicting individual choices, evaluated in the deterministic rule (accuracy) and the probabilistic rule (cross-entropy). The curves in Fig. 5 are smooth when approaching the right and left ends, which suggest the smooth convergence of TB-ResNets towards DCMs as δ decreases and towards DNNs as δ increases. The TB-ResNet framework can outperform DCMs because the DNN part in the TB-ResNets relaxes the stringent structure constraint embedded in DCMs, which nearly always leads to misspecification errors and underfitting. The TB-ResNet framework can outperform DNNs because of the localization in the DCM part and the regularization by the small δ . Overall, the results here clearly demonstrate the TB-ResNets can outperform both DCMs and DNNs, incorporate the two model families as two specific cases, and can converge to the two ends with δ approaching zero or one.

The optimum δ values, as highlighted in bold font in Table 3, suggest that the MNL and HD theories are more complete than the PT. As marked by the middle (3rd) model in each panel, the optimum δ values equal to 0.008 and 0.05 in the MNL and HD scenarios, while that in the PT scenario is around 0.9. As introduced in Section 3.3, an optimum and small δ strongly regularizes the DNN part, suggesting relative completeness of the DCM part; while an optimum and large δ uses the DNN part to search in a larger function space, suggesting that the DCM part is less complete. Therefore, our results imply that the MNL and HD theories are relatively complete in capturing the main behavioral mechanism, while the PT has more room for improvement. This finding can also be validated by the gap of the predictive performance between DCMs and DNNs in the three scenarios. A larger prediction gap is related to a larger optimum δ value: the optimum δ values are 0.008, 0.05, and 0.9 in MNL, HD, and PT scenarios, corresponding to the prediction accuracy gaps of 5.2% (=55.8%–50.6%), 14.5% (=71.2%–56.7%), and 19.1% (=88.3%–69.2%). Intuitively, a larger predictive gap between the DCMs and the DNNs suggests more incompleteness of the DCMs. Therefore, the empirical results highly align with our theoretical discussions. The empirical optimum δ value creates an optimum TB-ResNet model and functions as a diagnostic tool to evaluate the completeness of the theories.

Although Table 3 presents the optimum δ as exact values, the optimum δ is more likely to take a range of values, caused by the potential inconsistency of the multiple predictive metrics (Wang et al., 2020c). Model selections based on multiple predictive metrics can be inconsistent with each other. For example, in Table 3 and Fig. 5, while DNNs outperform PT in both prediction accuracy and cross-entropy loss, DNNs outperform MNL and HD models in only prediction accuracy but underperform in cross-entropy loss. In selecting the optimum TB-ResNets, the optimum δ based on the prediction accuracy and the cross-entropy loss is always slightly different. The authors would argue that this optimum range is actually more preferred to an exact optimum value, as it presents flexibility for further design of δ to enrich the TB-ResNet framework.

The TB-ResNets' predictive performance varies with the design of the DNN part. Table 3 presents the TB-ResNets constructed by two DNN parts with three and five layers, reported respectively in columns 2–4 and 5–7. The TB-ResNet with the three-layer DNN architecture slightly outperforms that with the five-layer DNN part in the MNL scenario, while the former underperforms the latter in the PT and HD scenarios. Therefore, it is possible to improve the TB-ResNets by adopting more effective DNN architecture, given that a vast number of available DNN architectures exist. However, this work seeks to demonstrate that, regardless of the specific DNN architecture, the TB-ResNets can always improve the performance of DNNs through the combination with the DCMs. Hence, in-depth exploration into either the DNN or the DCM part is beyond the scope of our work.

5.4. Robustness

Although robustness is not a traditional topic in travel demand modeling, it is increasingly important in both general ML discussions and the specific setting of demand modeling, since it can measure the local regularity of economic information. A model lacking robustness tends to predict an irregular behavioral pattern in which a slight perturbation of the input leads to a large change of the output, which is behaviorally unrealistic since a tiny price increase should not dramatically change

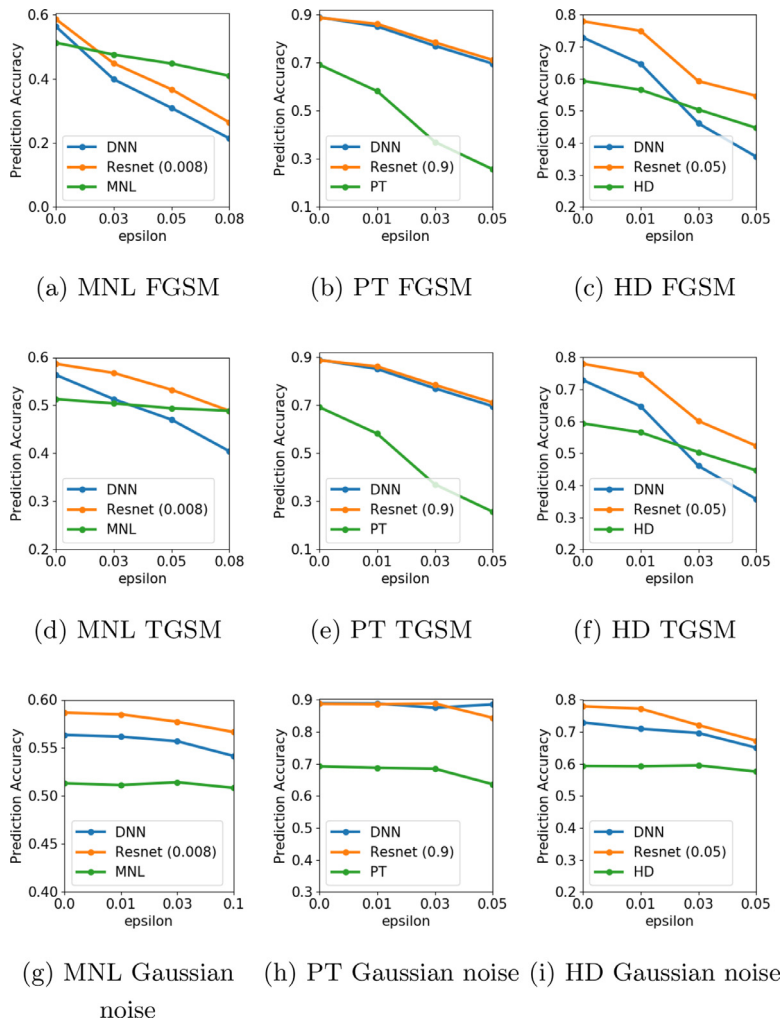


Fig. 6. Prediction accuracy with three perturbations (Gaussian noise, FGSM, and TGS) using the perturbed testing sets

decision-maker’s choice. In practice, perturbations can arise from the measurement noises in data or the randomness in data collection process.⁶

This type of behavioral regularity can be formally measured by robustness, or more specifically, the change of predictive performance under small perturbations of the inputs. This study uses three perturbations including two adversarial attacks (FGSM and TGS) and one random noise (Gaussian noise). They are (1) the fast gradient sign attack method (FGSM) ($x_{adv} = x + \epsilon \times \text{sign}(\nabla_x L(y, \hat{y}))$), and (2) the target gradient sign method (TGS) ($x_{adv} = x - \epsilon \times \text{sign}(\nabla_x L(y_{target}, \hat{y}))$) (Goodfellow et al., 2015; Kurakin et al., 2016), and (3) the Gaussian noise (GN) ($x_{adv} = x + \epsilon \times \Delta x$, where x is already standardized and $\Delta x \sim \mathcal{N}(0, 1)$). The models are trained on the initial data points x , perturbed with the x_{adv} under the three mechanisms, and evaluated by using the x_{adv} to predict the y corresponding to the initial x . A robust model should respond modestly to the perturbation, since the small input perturbations should lead to a small change of outputs. The Gaussian noise perturbation simulates random noises in data sets, while FGSM and TGS methods simulate malicious attacks on the system. The adversarial attacks are common methods to evaluate system robustness in the ML literature, while Gaussian noise is more realistic in approximating random noise in the context of choice modeling.

MNL-ResNets and HD-ResNets are more robust than DNNs under the two adversarial attacks, as shown in Figs. 6a, 6d, 6c, and 6f. First, DCMs are more robust than DNNs. For example, in Fig. 6a, while the prediction accuracy of the MNL model is lower than that of the DNN model when $\epsilon = 0$, the accuracy of the MNL model becomes much higher than that of the DNN when $\epsilon > 0.03$. Since the MNL-ResNet and the HD-ResNet models are the combination of DCMs and DNNs, the two TB-ResNets are more robust than DNNs, as shown by the orange curves lying above the blue curves in Figs. 6a, 6d, 6c,

⁶ For more details, readers can refer to Wang et al. (2020b) to understand the relationship between robustness and local regularity of economic information.

Table 4
Comparison of TB-ResNets to DCMs and DNNs.

Models	Prediction	Interpretability	Robustness
Compared to DNNs	Marginal improvement (by stabilization and regularization)	Significant improvement (by stabilization and regularization)	Significant improvement (by stabilization and regularization)
Compared to DCMs	Significant improvement (by augmenting and enriching utility function)	Significant improvement (by augmenting and enriching utility function)	No improvement

and 6f. The reason is that the DCMs are functioning as an anchoring skeleton in the TB-ResNet system, so the local utility patterns of TB-ResNets are less irregular than DNNs. This robustness intuition is closely related to the previous discussions about utility patterns: when utility patterns are more regular, the system tends to be more robust.

However, the robustness pattern under the PT scenario is quite different from MNL and HD, because of the particularity of PT and the different optimum δ value in PT-ResNets. While MNL and HD theories are designed to have smooth and well-bounded input gradients, the probability weighting function in the PT can have an input gradient that is close to infinity when the probabilities approach zero. The underlying behavioral intuition is that people tend to significantly exaggerate very small winning chances, such as in gambling, leading to a large overestimation of the utility gain associated with the alternative. In other words, PT is designed to be sensitive to perturbations, thus lacking robustness, particularly around the region of small probabilities. This can be seen in Figs. 6b and 6e: the PT models not only have smaller prediction accuracy than the DNNs as $\epsilon = 0$, but also decrease much more quickly than the DNNs. But on the other side, the prediction accuracy decrease of PT-ResNets and DNNs largely align with each other, because of the large optimum δ value in PT-ResNet. When δ is close to one, the PT-ResNet largely resembles the DNN part, so their robustness performance tends to be similar. Nonetheless, the pattern of an infinite gradient seems to be very specific to PT, because generally a DCM should have regular and bounded gradients, thus stabilizing the TB-ResNets rather than destabilizing them.

Under the Gaussian noise, changes in the predictive performance are much more modest, although the pattern still appears similar to that from the adversarial attacks. As shown in Figs. 6g, 6h, and 6i, all of the curves are nearly flat with slight downward sloping, suggesting that the prediction accuracy does not vary much under Gaussian noise. The difference between Gaussian noise and adversarial attacks is no surprise, since the adversarial attacks target the most vulnerable part of the input space while Gaussian noise is random. However, even with the relatively flat curves, the findings here appear similar to those from the adversarial attacks. In the MNL and HD scenarios, MNL and HD models are slightly more robust than the DNNs, while the PT model seems slightly less so than the DNNs. As a result, the MNL-ResNet and the HD-ResNet appear slightly more robust than the DNNs in the MNL and HD scenarios, while the PT-ResNet and the DNN in the PT scenario nearly overlap.

6. Conclusions and Discussions

This study introduces a TB-ResNet framework to analyze individual decision-making by synergizing the theory- and data-driven methods, based on the utility interpretation shared by DNNs and DCMs. The TB-ResNet framework can be understood from the perspectives of architecture design, model ensemble, gradient boosting, regularization, function approximation, and theory diagnosis. Three instances of TB-ResNets, including MNL-ResNets, PT-ResNets and HD-ResNets, are created and tested empirically on three data sets to evaluate model prediction, interpretability, and robustness.

As summarized in Table 4, our empirical results demonstrate that TB-ResNets are more predictive, interpretable, and robust overall than pure DCMs and DNNs, although several exceptions exist. Compared to DNNs, TB-ResNets are more predictive, interpretable, and robust because the DCM component in TB-ResNets can stabilize the utility functions and regularize the DNN component. Compared to DCMs, TB-ResNets are more predictive and interpretable because richer utility functions are augmented to the skeleton DCM by the DNN component in TB-ResNets. The TB-ResNets are formulated with a flexible $(\delta, 1 - \delta)$ weighting, thus taking advantage of both the simplicity of the DCMs and the richness of the DNNs, preventing the underfitting of the DCMs and the overfitting of the DNNs, and providing insights into the completeness of the DCM theories. Our findings are consistent across the three scenarios (MNL, PT, and HD). While exceptions exist in the PT scenario and the comparison to DCMs for robustness evaluation, our main findings hold from both theoretical and empirical perspectives.

The TB-ResNet is a method to reconcile the handcrafted and the automatic utility specification, which resonates with the broad discussions about how to interact the general-purpose ML and domain-specific models. DNNs generate dominant predictive performance because they can automatically learn utility specification, termed as an “end-to-end” system that can “learn from scratch”. This power of automation is treated as its main strength over traditional methods that rely on domain knowledge to handcraft utility functions (Mullainathan and Spiess, 2017; LeCun et al., 2015). However on the other side, researchers argued that automatic feature learning with zero prior knowledge does not appear to be a viable approach. Liao and Poggio (2018) contended that “being lazy (automatic learning) is good, but being too lazy is not”. Our TB-ResNet is a tangible example of integrating handcrafted and automatic learning systems in the context of individual decision-making.

This TB-ResNet framework reveals the particularity of DNNs among all the ML classifiers, since DNNs are closely related to the DCMs through the implicit utility interpretation and the probabilistic behavioral perspective. DNNs can model the

choice probabilities using the Softmax activation function, and are thus highly compatible with the classical probabilistic behavioral modeling. On the contrary, many other ML classifiers adopt a deterministic approach in modeling the choice outputs, such as K nearest neighbors and support vector machines, and are therefore somewhat more difficult to synergize with DCMs. The TB-ResNets take advantage of the utility interpretation for this synergy, so this approach resonates better with the classical utility theory than the perspectives of model ensemble and gradient boosting. From the architecture design perspective, the TB-ResNet framework can be further improved by incorporating a large number of DNN architectures from the ML community (He et al., 2016; Krizhevsky et al., 2012) or designing new DNN architectures with behavioral knowledge (Wang et al., 2020a). The architecture design perspective will become even more important when modelers start to use high-dimensional inputs such as imagery and natural language for demand modeling.

Overall, this study introduces a simple but flexible TB-ResNet framework. It is simple because it combines DCMs and DNNs neatly into the TB-ResNet architecture, and is flexible because it can incorporate any DCM as the theory-driven part and any DNN as the data-driven part. The three instances of the TB-ResNets provide evidence that the TB-ResNet framework is malleable enough to be applied to diverse decision scenarios with many benefits. Despite its strength, many questions still remain. Although the combination of the DNNs and DCMs can generate mutual benefits, the TB-ResNet is not the only possible approach to improve prediction, interpretation, and robustness, simply because a large number of other methods have been developed to improve DCMs and DNNs in each community. DCMs can be more predictive with richer utility specifications, and DNNs can be more predictive with better designs in architecture, hyperparameters, and training algorithms. DNNs can be made more interpretable and robust by adopting activation maximization (AM) (Erhan et al., 2009), LIME (Ribeiro et al., 2016), and minimax training procedures (Madry et al., 2017). Future studies can further enrich our TB-ResNet framework by using other DCMs for the DCM component and other DNN architectures for the DNN component. This study highlights a synergetic perspective and a comprehensive model evaluation based on three criteria, so future studies could take the complementarity of the data-driven and theory-driven methods beyond simple prediction comparison. We hope that this work can pave the way for future studies to create more links between the data-driven and theory-driven methods, because their complementary nature provides immense opportunities, their underlying perspectives are interwoven, and their synergy can overcome their respective weaknesses.

Acknowledgement

The research is supported by the National Research Foundation (NRF), Prime Ministers Office, Singapore, under CREATE programme, Singapore-MIT Alliance for Research and Technology (SMART) Centre, Future Urban Mobility (FM) IRG. We thank Nick Caros for his careful proofreading.

Appendix A. Proof of Propositions 1 and 2

Proposition 1 Proof. This proof can be found in many textbooks (Train 2009; Ben-Akiva and Lerman 1985). With Gumbel distributional assumption, Equation 3 could be solved in an analytical way:

$$\begin{aligned}
 P_{ik} &= \int_{-\infty}^{+\infty} \prod_{j \neq k} e^{-e^{-(V_{ik}-V_{ij}+\epsilon_{ik})}} f(\epsilon_{ik}) d\epsilon_{ik} \\
 &= \int \prod_j e^{-e^{-(V_{ik}-V_{ij}+\epsilon_{ik})}} e^{-\epsilon_{ik}} d\epsilon_{ik} \\
 &= \int \exp(e^{-\epsilon_{ik}} \sum_j -e^{-(V_{ik}-V_{ij})}) e^{-\epsilon_{ik}} d\epsilon_{ik} \\
 &= \int_{-\infty}^0 \exp(-t \sum_j e^{-(V_{ik}-V_{ij})}) dt \\
 &= \frac{e^{V_{ik}}}{\sum_j e^{V_{ij}}}
 \end{aligned} \tag{36}$$

in which the fourth equation uses $t = e^{-\epsilon_{ik}}$. Note this formula in Equation 36 is the Softmax function in DNN. V_{ik} is both the deterministic utility in RUM and the inputs into the Softmax function in DNN.

Proposition 2 Proof. The detailed proof of Proposition 2 could be found in Lemma 2 of McFadden’s seminal paper (McFadden 1974). Here is a brief summary. Suppose that one individual i firstly chooses between alternative k and T alternatives j . Then according to Equations 3 and 36,

$$\begin{aligned}
 P_{ik} &= \frac{e^{V_{ik}}}{e^{V_{ik}} + T e^{V_{ij}}} \\
 &= \int F(\epsilon_{ik} + V_{ik} - V_{ij})^T dF(\epsilon_{ik})
 \end{aligned} \tag{37}$$

Suppose that the individual i chooses between alternatives k and alternative l in another choice scenario, and alternative l is constructed such that $Te^{V_{ij}} = e^{V_{il}}$. Then

$$\begin{aligned}
 P_{ik} &= \frac{e^{V_{ik}}}{e^{V_{ik}} + e^{V_{il}}} \\
 &= \int F(\epsilon_{ik} + V_{ik} - V_{il})dF(\epsilon_{ik}) \\
 &= \int F(\epsilon_{ik} + V_{ik} - V_{ij} - \log T)dF(\epsilon_{ik})
 \end{aligned}
 \tag{38}$$

By construction, Equations 37 and 38 are equivalent

$$\int F(\epsilon_{ik} + V_{ik} - V_{ij} - \log T) - F(\epsilon_{ik} + V_{ik} - V_{ij})^T dF(\epsilon_{ik}) = 0$$

Since $F(\epsilon)$ is transition complete, meaning that $\forall a, Eh(\epsilon + a) = 0$ implies $h(\epsilon) = 0, \forall \epsilon$, it implies

$$F(V_{ik} - \log T) = F(V_{ik})^T, \forall V_{ik}, T$$

Taking $V_{ik} = 0$ implies $F(-\log T) = e^{-\alpha T}$. Taking $V_{ik} = \log T - \log L$ implies $F(-\log L) = F(\log T/L)^T$. Hence $F(\log T/L) = F(-\log L)^{1/T} = e^{-\alpha L/T}$. Therefore, $F(\epsilon) = e^{-\alpha e^{-\epsilon}}$. This is the function of Gumbel distribution when $\alpha = 1$.

Appendix B. Summary Statistics of Three Data Sets

Table 5
Summary Statistics of SG data set

Variables			Variables		
Name	Mean	Std.	Name	Mean	Std.
Male (Yes = 1)	0.383	0.486	Age <35 (Yes = 1)	0.329	0.470
Age>60 (Yes = 1)	0.075	0.263	Low education (Yes = 1)	0.331	0.471
High education (Yes = 1)	0.480	0.500	Low income (Yes = 1)	0.035	0.184
High income (Yes = 1)	0.606	0.489	Full job (Yes = 1)	0.602	0.490
Walk: walk time (min)	60.50	54.88	Bus: cost (\$SG)	2.070	1.266
Bus: walk time (min)	11.96	10.78	Bus: waiting time (min)	7.732	5.033
Bus: in-vehilce time (min)	25.06	18.91	RideSharing: cost (\$SG)	14.48	11.64
RideSharing: waiting time (min)	7.108	4.803	RideSharing: in-vehilce time (min)	18.28	13.39
AV: cost (\$SG)	16.08	14.60	AV: waiting time (min)	7.249	5.674
AV: in-vehilce time (min)	20.11	16.99	Drive: cost (\$SG)	10.49	10.57
Drive: walk time (min)	3.968	4.176	Drive: in-vehilce time (min)	17.43	14.10
Statistics					
Number of samples	8418				
Number of choices	Walk: 874 (10.38%); Bus: 1951 (23.18%); RideSharing: 904 (10.74%); Drive 3774 (44.83%); AV: 915 (10.87%)				

Table 6
Summary Statistics of PT data set

Variables			Variables		
Name	Mean	Std.	Name	Mean	Std.
Male (Yes = 1)	0.619	0.485	Age	47.46	12.89
Num of years in school	6.746	3.821	Household annual income (1 million dong)	20.27	21.15
Chinese (Yes = 1)	0.055	0.228	Distance to the nearest local market (km)	1.482	1.840
Living in Southern Vietnam (Yes = 1)	0.541	0.498	Reward 1 in option A (1 million dong)	0.032	0.016
Prob of reward 1 in option A	0.638	0.263	Reward 2 in option A (1 million dong)	0.016	0.015
Prob of reward 2 in option A	0.362	0.263	Reward 1 in option B (1 million dong)	0.076	0.038
Prob of reward 1 in option B	0.486	0.252	Reward 2 in option B (1,000 dong)	-0.340	9.640
Prob of reward 2 in option B	0.514	0.252			
Statistics					
Number of samples	5249				
Number of choices	Option A: 2823 (53.78%); Option B: 2426 (46.22%)				

Table 7
Summary Statistics of HD data set

Variables			Variables		
Name	Mean	Std.	Name	Mean	Std.
Male (Yes = 1)	0.618	0.486	Age	47.51	12.94
Num of years in school	6.764	3.843	Household income (1 million dong)	20.71	21.23
Chinese (Yes = 1)	0.055	0.228	Distance to the nearest local market (km)	1.506	1.846
Living in Southern Vietnam (Yes = 1)	0.534	0.499	^a Trusted agent (Yes = 1)	0.028	0.165
Payment received in the risk experiment (1 million dong)	20.97	21.17	Amount of immediate reward (1 million dong)	0.075	0.078
Amount of delayed reward (1 million dong)	0.150	0.104	Days of delay	35.67	32.33
Statistics					
Number of samples	5340				
Number of choices	Immediate reward: 2670 (50.0%); Future reward: 2670 (50.0%)				

^a Trusted agents are people who would keep the money until delayed delivery date to ensure subjects believed the money would be delivered. The selected trusted persons were usually village heads or presidents of womens associations.

Appendix C. Results of Simultaneous Training

Table 8
Performance of DCMs, DNNs, and TB-ResNets in testing sets (Sequential and simultaneous training).

	Prediction accuracy (Sequential)	Cross-entropy loss (Sequential)	F1 score (Sequential)	Prediction accuracy (Simultaneous)	Cross-entropy Loss (Simultaneous)	F1 score (Simultaneous)	Baseline (Largest share)
Panel 1. Performance of MNL models							
MNL	50.6%	1.254	0.439	50.6%	1.254	0.439	44.8%
MNL ResNet ($\delta = 1e-5$)	53.1%	1.207	0.485	52.1%	1.224	0.468	44.8%
MNL ResNet ($\delta = 0.008$)	57.0%	1.237	0.559	56.6%	1.150	0.542	44.8%
MNL ResNet ($\delta = 0.05$)	56.1%	1.572	0.559	56.1%	1.213	0.550	44.8%
DNN	55.8%	2.861	0.555	55.8%	2.861	0.555	44.8%
Panel 2. Performance of PT models							
PT	69.2%	0.602	0.693	69.2%	0.602	0.693	53.8%
PT ResNet ($\delta = 1e-05$)	75.6%	0.502	0.762	76.7%	0.477	0.772	53.8%
PT ResNet ($\delta = 0.9$)	89.2%	0.343	0.887	88.7%	0.347	0.882	53.8%
PT ResNet ($\delta = 0.99$)	88.6%	0.318	0.882	88.6%	0.335	0.884	53.8%
DNN	88.3%	0.353	0.885	88.3%	0.353	0.885	53.8%
Panel 3. Performance of HD models							
HD	56.7%	0.684	0.568	56.7%	0.684	0.568	50.0%
HD ResNet ($\delta = 1e-05$)	68.9%	0.523	0.689	68.7%	0.517	0.686	50.0%
HD ResNet ($\delta = 0.05$)	77.6%	0.437	0.764	77.3%	0.439	0.774	50.0%
HD ResNet ($\delta = 0.99$)	76.0%	0.444	0.763	76.3%	0.440	0.774	50.0%
DNN	71.2%	0.909	0.722	71.2%	0.909	0.722	50.0%

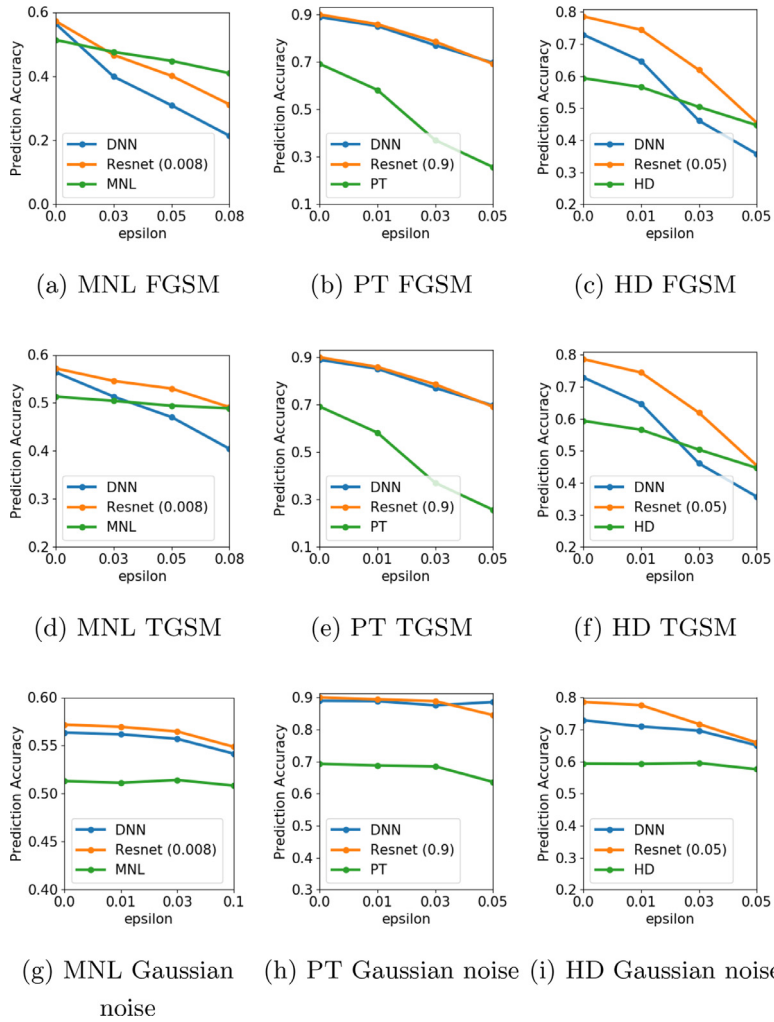


Fig. 7. Prediction accuracy with perturbations (Gaussian noise, FGSM, and TGSM attacks, simultaneous training)

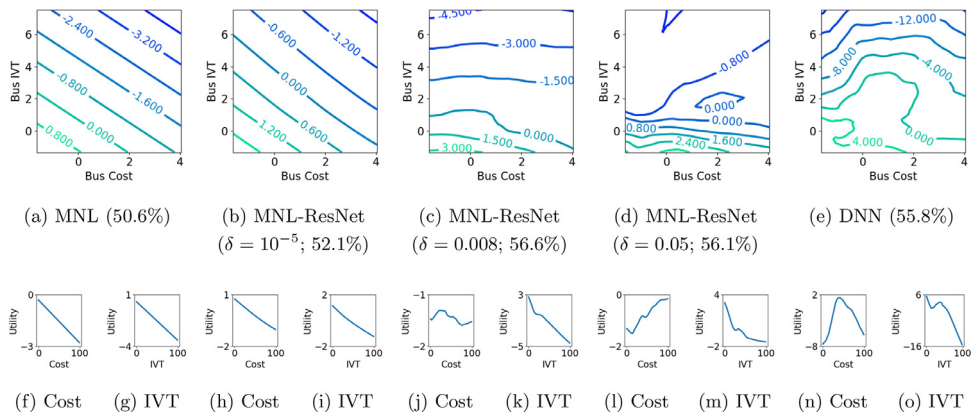


Fig. 8. Utility functions of MNL-ResNets, MNL, and DNNs (Simultaneous training).

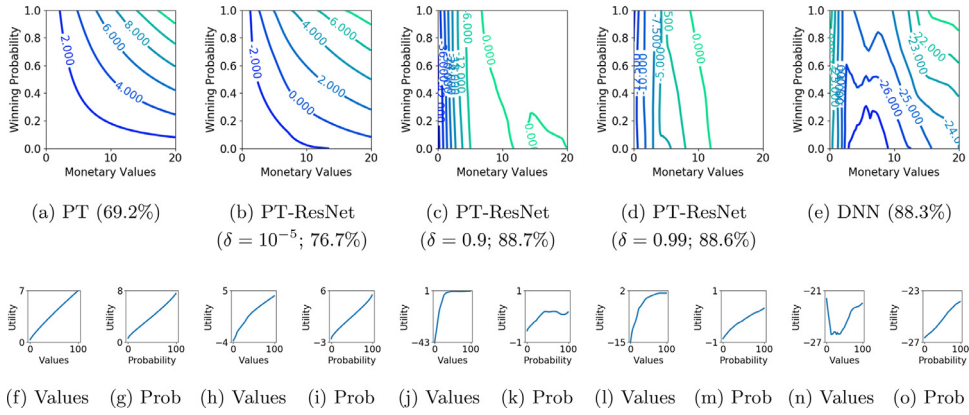


Fig. 9. Utility functions of PT-ResNets, PT, and DNNs (Simultaneous training).

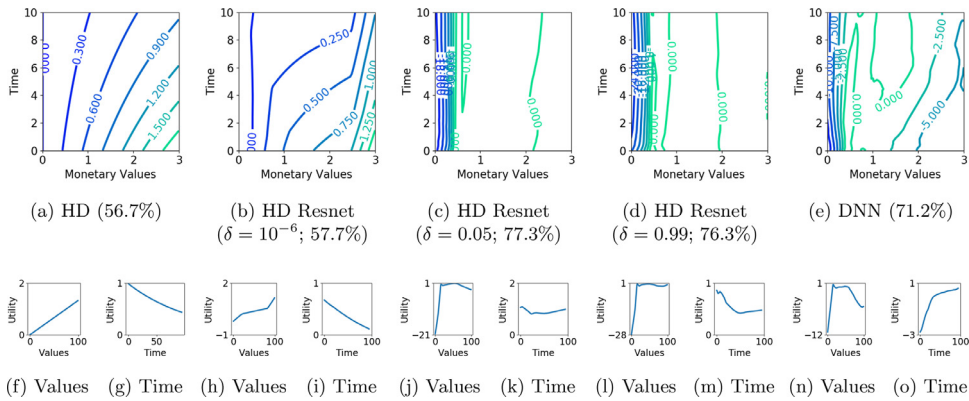


Fig. 10. Utility functions of HD-ResNets, HD, and DNNs (Simultaneous training)

Appendix D. Proof of Proposition 4

Proof. By using the definition of Rademacher complexity and Proposition 3, the left hand side can be rewritten as:

$$\mathbb{E}_S[L(\hat{f}) - L(f^*_f)] \leq 2\mathbb{E}_S\hat{\mathcal{R}}_n(\mathcal{F}|_S) \tag{39}$$

$$= 2\mathbb{E}_S\hat{\mathcal{R}}_n((1 - \delta)\mathcal{F}_1 + \delta\mathcal{F}_2|_S) \tag{40}$$

$$= 2\mathbb{E}_S\mathbb{E}_\epsilon \sup_{f_1 \in \mathcal{F}_1; f_2 \in \mathcal{F}_2} \frac{1}{N} \langle \epsilon, (1 - \delta)f_1 + \delta f_2 \rangle \tag{41}$$

$$\leq 2\mathbb{E}_S \left[\mathbb{E}_\epsilon \sup_{f_1 \in \mathcal{F}_1} \frac{1}{N} \langle \epsilon, (1 - \delta)f_1 \rangle + \mathbb{E}_\epsilon \sup_{f_2 \in \mathcal{F}_2} \frac{1}{N} \langle \epsilon, \delta f_2 \rangle \right] \tag{42}$$

$$= 2\mathbb{E}_S[(1 - \delta)\hat{\mathcal{R}}_n(\mathcal{F}_1|_S) + \delta\hat{\mathcal{R}}_n(\mathcal{F}_2|_S)] \tag{43}$$

in which the first line uses Proposition 3; the second line uses the definition of \mathcal{F} ; the third line uses a definition $f := (1 - \delta)f_1 + \delta f_2$; the fourth line uses the convexity of the sup operator; the last one uses the definition of Rademacher complexity again.

Appendix E. Utility Intuition for Five-Layer DNN Architecture

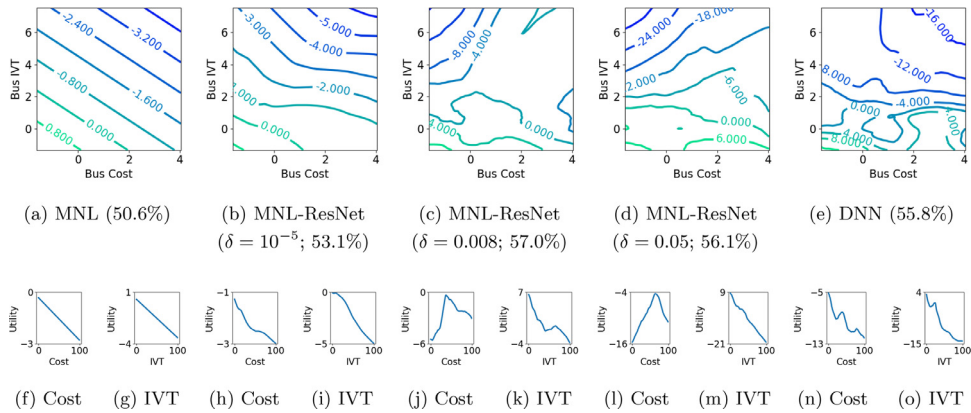


Fig. 11. Utility functions of MNL-ResNets, MNL, and DNNs; upper row: visualization of 2D utility functions; lower row: visualization of 1D utility functions; percentages on the upper row are the prediction accuracy of the five models.

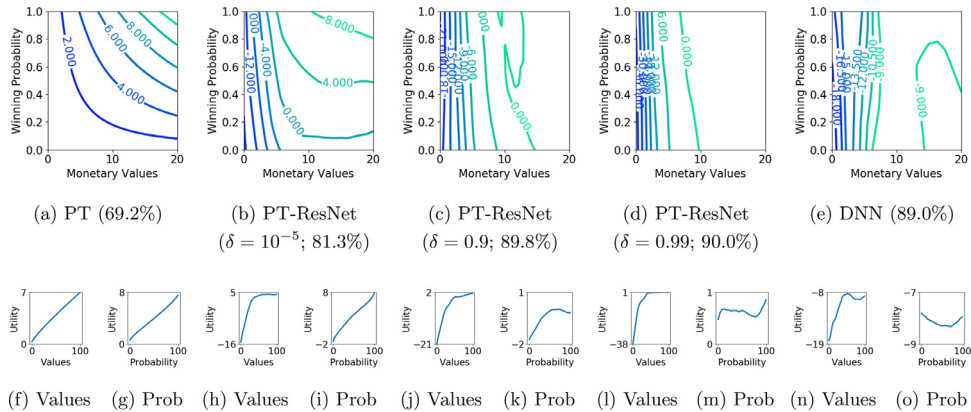


Fig. 12. Utility functions of PT-ResNets, PT, and DNNs; upper row: visualization of 2D utility functions; lower row: visualization of 1D utility functions; percentages on the upper row are the prediction accuracy of the five models.

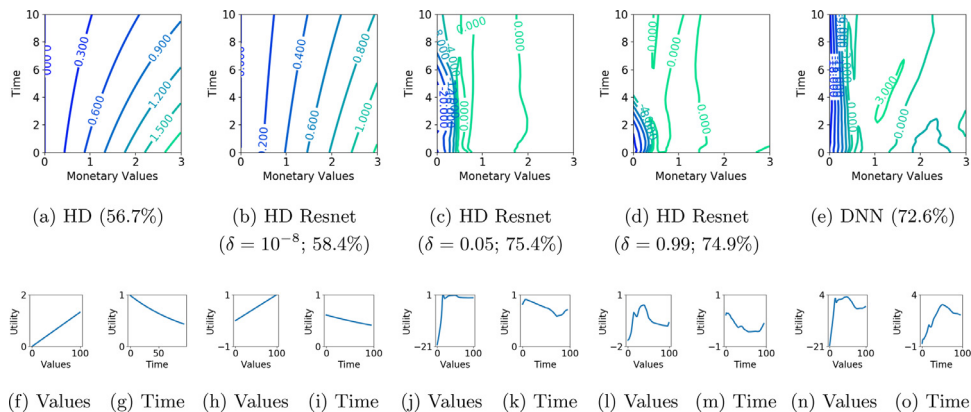


Fig. 13. Utility functions of HD-ResNets, HD, and DNNs; upper row: visualization of 2D utility functions; lower row: visualization of 1D utility functions; percentages on the upper row are the prediction accuracy of the five models.

Appendix F. Convergence of Model Training


Table 9 summarizes the model performance averaged over the last 200 epoches. There exists only a minor performance difference between Table 9 and Table 3, suggesting that our findings in Table 3 are robust. Take the 3-layer MNL-ResNet ($\delta = 0.008$) as an example. The prediction accuracy, cross-entropy loss, and F1 score in Table 3 are 57.0%, 1.237, and 0.559, as opposed to 57.0%, 1.236, and 0.559 in Table 9. Other numbers of the two tables are also similar.

Fig. 14 visualizes the convergence of the performance in both training and testing sets, with the green and the blue curves representing the prediction errors and the cross-entropy losses. As shown by Figs. 5b and 5c, both PT-ResNet and HD-ResNet have converged after around 4000 epoches, although the MNL-ResNet demonstrates a slight overfitting measured by the testing loss, as shown in Fig. 5a. The slight overfitting is caused by the potential inconsistency between the two evaluation metrics (as in Fig. 5), which is beyond the scope of this work. It is also possible to further improve the TB-ResNets' performance by choosing the stopping point more strategically using a separate validation set. Nonetheless, the three figures have shown that the training and testing performances are largely stable in the last around 1,000 epoches across three TB-ResNets, consistent with our finding in Table 9.

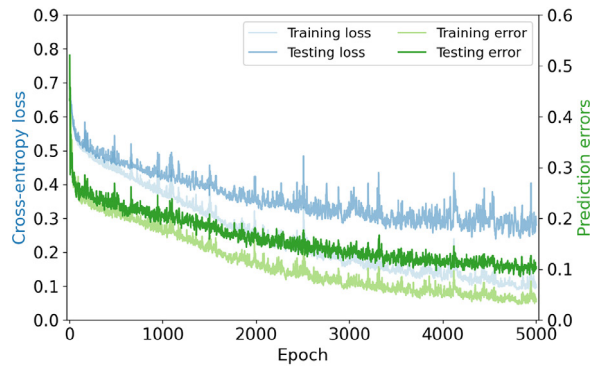
Table 9

Average performance of the last 200 epoches in DCMs, DNNs, and TB-ResNets in testing sets (Sequential training).


	Prediction accuracy (3-layer)	Cross-entropy loss (3-layer)	F1 score (3-layer)	Prediction accuracy (5-layer)	Cross-entropy Loss (5-layer)	F1 score (5-layer)	Baseline (Largest share)
Panel 1. Performance of MNL models							
MNL	50.6%	1.255	0.439	50.6%	1.255	0.439	44.8%
MNL ResNet ($\delta = 1e-5$)	53.1%	1.207	0.485	53.1%	1.346	0.498	44.8%
MNL ResNet ($\delta = 0.008$)	57.0%	1.236	0.559	55.6%	3.214	0.557	44.8%
MNL ResNet ($\delta = 0.05$)	56.5%	1.552	0.559	57.1%	3.338	0.576	44.8%
DNN	55.7%	2.824	0.555	55.5%	4.073	0.548	44.8%
Panel 2. Performance of PT models							
PT	69.2%	0.603	0.693	69.2%	0.602	0.693	53.8%
PT ResNet ($\delta = 1e-05$)	76.0%	0.503	0.762	81.1%	0.433	0.813	53.8%
PT ResNet ($\delta = 0.9$)	89.2%	0.349	0.887	90.1%	0.340	0.894	53.8%
PT ResNet ($\delta = 0.99$)	88.2%	0.327	0.882	89.9%	0.340	0.899	53.8%
DNN	88.0%	0.353	0.880	89.5%	0.360	0.898	53.8%
Panel 3. Performance of HD models							
HD	57.7%	0.681	0.578	57.7%	0.681	0.578	50.0%
HD ResNet ($\delta = 1e-05$)	66.9%	0.532	0.657	73.6%	0.471	0.737	50.0%
HD ResNet ($\delta = 0.05$)	77.1%	0.439	0.764	75.3%	0.454	0.761	50.0%
HD ResNet ($\delta = 0.99$)	76.7%	0.436	0.763	75.8 %	0.448	0.760	50.0%
DNN	73.1%	0.853	0.709	72.4%	1.212	0.726	50.0%



(a) MNL-ResNet



(b) PT-ResNet



(c) HD-ResNet

Fig. 14. Performance of TB-ResNets over the training process

References

- Arrow, K.J., 1965. Aspects of the theory of risk-bearing. Yrjo Jahnssonin Saatio.
- Avineri, E., Bovy, P., 2008. Identification of parameters for a prospect theory model for travel choice analysis. *Transportation Research Record: Journal of the Transportation Research Board* (2082) 141–147.
- Baehrens, D., Schroeter, T., Harmeling, S., Kawanabe, M., Hansen, K., Mller, K.-R., 2010. How to explain individual classification decisions. *Journal of Machine Learning Research* 11 (Jun), 1803–1831.
- Ben-Akiva, M.E., Lerman, S.R., 1985. *Discrete choice analysis: theory and application to travel demand*, 9. MIT press.
- Bentz, Y., Merunka, D., 2000. Neural networks and the multinomial logit for brand choice modelling: a hybrid approach. *Journal of Forecasting* 19 (3), 177–200.
- Borysov, S.S., Rich, J., Pereira, F.C., 2019. How to generate micro-agents? a deep generative modeling approach to population synthesis. *Transportation Research Part C: Emerging Technologies* 106, 73–97.
- Bradley, M.A., Daly, A.J., 1997. Estimation of logit choice models using mixed stated preference and revealed preference information. *Understanding travel behaviour in an era of change* 209–232.
- Camerer, C.F., Kunreuther, H., 1989. Decision processes for low probability events: Policy implications. *Journal of Policy Analysis and Management* 8 (4), 565–592.
- Cantarella, G.E., de Luca, S., 2005. Multilayer feedforward networks for transportation mode choice analysis: An analysis and a comparison with random utility models. *Transportation Research Part C: Emerging Technologies* 13 (2), 121–155.
- van Cranenburgh, S., Alwosheel, A., 2019. An artificial neural network based approach to investigate travellers decision rules. *Transportation Research Part C: Emerging Technologies* 98, 152–166.
- Cybenko, G., 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* 2 (4), 303–314.
- De Palma, A., Ben-Akiva, M., Brownstone, D., Holt, C., Magnac, T., McFadden, D., Moffatt, P., Picard, N., Train, K., Wakker, P., 2008. Risk, uncertainty and discrete choice models. *Marketing Letters* 19 (3–4), 269–285.
- Dhami, S., 2016. *The Foundations of Behavioral Economic Analysis*. Oxford University Press.
- Do, L.N., Vu, H.L., Vo, B.Q., Liu, Z., Phung, D., 2019. An effective spatial-temporal attention based neural network for traffic flow prediction. *Transportation research part C: emerging technologies* 108, 12–28.
- Doshi-Velez, F., Kim, B., 2017. *Towards a rigorous science of interpretable machine learning*.
- Duan, Y., Lv, Y., Liu, Y.-L., Wang, F.-Y., 2016. An efficient realization of deep learning for traffic data imputation. *Transportation research part C: emerging technologies* 72, 168–181.
- Erhan, D., Bengio, Y., Courville, A., Vincent, P., 2009. Visualizing higher-layer features of a deep network. *University of Montreal* 1341 (3), 1.
- Freitas, A.A., 2014. *Comprehensible classification models: a position paper*. *ACM SIGKDD explorations newsletter* 15 (1), 1–10.
- Friedman, J.H., 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* 1189–1232.
- Glaeser, E.L., Kominers, S.D., Luca, M., Naik, N., 2018. Big data and big cities: The promises and limitations of improved measures of urban life. *Economic Inquiry* 56 (1), 114–137.
- Golowich, N., Rakhlin, A., Shamir, O., 2017. Size-independent sample complexity of neural networks. *arXiv preprint arXiv:1712.06541*.
- Goodfellow, I.J., Shlens, J., Szegedy, C., 2015. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Hagenauer, J., Helbich, M., 2017. A comparative study of machine learning classifiers for modeling travel mode choice. *Expert Systems with Applications* 78, 273–282.
- Hao, S., Lee, D.-H., Zhao, D., 2019. Sequence to sequence learning with attention mechanism for short-term passenger flow prediction in large-scale metro system. *Transportation Research Part C: Emerging Technologies* 107, 287–300.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hensher, D.A., Bradley, M., 1993. Using stated response choice data to enrich revealed preference discrete choice models. *Marketing Letters* 4 (2), 139–151.
- Hinton, G., Vinyals, O., Dean, J., 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hornik, K., 1991. Approximation capabilities of multilayer feedforward networks. *Neural networks* 4 (2), 251–257.
- Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2 (5), 359–366.
- Huang, X., Sun, J., Sun, J., 2018. A car-following model considering asymmetric driving behavior based on long short-term memory neural networks. *Transportation Research Part C: Emerging Technologies* 95, 346–362.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* 263–291.
- Karlaftis, M.G., Vlahogianni, E.I., 2011. Statistical methods versus neural networks in transportation research: Differences, similarities and some insights. *Transportation Research Part C: Emerging Technologies* 19 (3), 387–399.
- Kaur, S., Kremer, M., Mullainathan, S., 2015. Self-control at work. *Journal of Political Economy* 123 (6), 1227–1277.
- Koszegi, B., Rabin, M., 2006. A model of reference dependent preferences. *The Quarterly Journal of Economics* 1133–1165.
- Kotsiantis, S.B., Zaharakis, I., Pintelas, P., 2007. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering* 160, 3–24.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*, pp. 1097–1105.
- Kurakin, A., Goodfellow, I., Bengio, S., 2016. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*.
- Kurakin, A., Goodfellow, I., Bengio, S., 2017. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521 (7553), 436–444.
- Lee, S., Xie, K., Ngoduy, D., Keyvan-Ekbatani, M., 2019. An advanced deep learning approach to real-time estimation of lane-based queue lengths at a signalized junction. *Transportation research part C: emerging technologies* 109, 117–136.
- Liao, Q., Poggio, T., 2018. *When Is Handcrafting Not a Curse?* Technical Report.
- Lipton, Z.C., 2016. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*.
- Liu, E.M., 2013. Time to change what to sow: Risk preferences and technology adoption decisions of cotton farmers in china. *Review of Economics and Statistics* 95 (4), 1386–1403.
- Liu, L., Chen, R.-C., 2017. A novel passenger flow prediction model using deep learning methods. *Transportation Research Part C: Emerging Technologies* 84, 74–91.
- Loewenstein, G., Prelec, D., 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics* 107 (2), 573–597.
- Ma, T., Antoniou, C., Toledo, T., 2020. Hybrid machine learning algorithm and statistical time series model for network-wide traffic forecast. *Transportation Research Part C: Emerging Technologies* 111, 352–372.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A., 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- McFadden, D., 1974. *Conditional logit analysis of qualitative choice behavior*.
- Montavon, G., Samek, W., Muller, K.-R., 2018. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing* 73, 1–15.
- Mozolin, M., Thill, J.-C., Usery, E.L., 2000. Trip distribution forecasting with multilayer perceptron neural networks: A critical evaluation. *Transportation Research Part B: Methodological* 34 (1), 53–73.
- Mullainathan, S., Spiess, J., 2017. Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31 (2), 87–106.
- Neumann, J. v., Morgenstern, O., 1944. *Theory of games and economic behavior*.

- Nicholson, W., Snyder, C., 2012. Uncertainty and strategy. *Microeconomics Theory: Basic Principles and Extensions*.
- O'Donoghue, T., Rabin, M., 1999. Doing it now or later. *American Economic Review* 103–124.
- O'Donoghue, T., Rabin, M., 2001. Choice and procrastination. *The Quarterly Journal of Economics* 116 (1), 121–160.
- Omrani, H., 2015. Predicting travel mode of individuals by machine learning. *Transportation Research Procedia* 10, 840–849.
- Papernot, N., McDaniel, P., Goodfellow, I., 2016. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*.
- Papernot, N., McDaniel, P., Wu, X., Jha, S., Swami, A., 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In: 2016 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 582–597.
- Paredes, M., Hemberg, E., O'Reilly, U.-M., Zegras, C., 2017. Machine learning or discrete choice models for car ownership demand estimation and prediction? In: *Models and Technologies for Intelligent Transportation Systems (MT-ITS), 2017 5th IEEE International Conference on*. IEEE, pp. 780–785.
- Polson, N.G., Sokolov, V.O., 2017. Deep learning for short-term traffic flow prediction. *Transportation Research Part C: Emerging Technologies* 79, 1–17.
- Pratt, J.W., 1964. Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society* 122–136.
- Pulugurta, S., Arun, A., Errampalli, M., 2013. Use of artificial intelligence for mode choice analysis and comparison with traditional multinomial logit model. *Procedia-Social and Behavioral Sciences* 104, 583–592.
- Ribeiro, M.T., Singh, S., Guestrin, C., 2016. Why should i trust you?: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 1135–1144.
- Samuelson, P.A., 1937. A note on measurement of utility. *The Review of Economic Studies* 4 (2), 155–161.
- Sekhar, C.R., Madhu, E., 2016. Mode choice analysis using random forest decision trees. *Transportation Research Procedia* 17, 644–652.
- Seo, T., Kusakabe, T., Gotoh, H., Asakura, Y., 2017. Interactive online machine learning approach for activity-travel survey. *Transportation Research Part B: Methodological*.
- Sydner, J., 2010. (over) insuring modest risks. *American Economic Journal: Applied Economics* 2 (4), 177–199.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2014. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tanaka, T., Camerer, C.F., Nguyen, Q., 2010. Risk and time preferences: linking experimental and household survey data from vietnam. *American Economic Review* 100 (1), 557–571.
- Train, K., 1980. A structured logit model of auto ownership and mode choice. *The Review of Economic Studies* 47 (2), 357–370.
- Train, K.E., 2009. *Discrete choice methods with simulation*. Cambridge university press.
- Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5 (4), 297–323.
- Wainwright, M.J., 2019. *High-dimensional statistics: A non-asymptotic viewpoint*, 48. Cambridge University Press.
- Wang, S., Mo, B., Zhao, J., 2020. Deep neural networks for choice analysis: Architecture design with alternative-specific utility functions. *Transportation Research Part C: Emerging Technologies* 112, 234–251.
- Wang, S., Wang, Q., Bailey, N., Zhao, J., 2018. Deep neural networks for choice analysis: A statistical learning theory perspective. *arXiv preprint arXiv:1810.10465*.
- Wang, S., Wang, Q., Zhao, J., 2020. Deep neural networks for choice analysis: Extracting complete economic information for interpretation. *Transportation Research Part C: Emerging Technologies* 118, 102701.
- Wang, S., Wang, Q., Zhao, J., 2020. Multitask learning deep neural networks to combine revealed and stated preference data. *Journal of Choice Modelling* 100236.
- Wang, S., Zhao, J., 2019. Risk preference and adoption of autonomous vehicles. *Transportation Research Part A: Policy and Practice* 126, 215–229.
- Wu, X., Guo, J., Xian, K., Zhou, X., 2018. Hierarchical travel demand estimation using multiple data sources: A forward and backward propagation algorithmic framework on a layered computational graph. *Transportation Research Part C: Emerging Technologies* 96, 321–346.
- Wu, Y., Tan, H., Qin, L., Ran, B., Jiang, Z., 2018. A hybrid deep learning based traffic flow prediction method and its understanding. *Transportation Research Part C: Emerging Technologies* 90, 166–180.
- Xiao, G., Juan, Z., Zhang, C., 2016. Detecting trip purposes from smartphone-based travel surveys with artificial neural networks and particle swarm optimization. *Transportation Research Part C: Emerging Technologies* 71, 447–463.
- Yang, S., Ma, W., Pi, X., Qian, S., 2019. A deep learning approach to real-time parking occupancy prediction in transportation networks incorporating multiple spatio-temporal data sources. *Transportation Research Part C: Emerging Technologies* 107, 248–265.
- Zhang, J., Zheng, Y., Qi, D., Li, R., Yi, X., Li, T., 2018. Predicting citywide crowd flows using deep spatio-temporal residual networks. *Artificial Intelligence* 259, 147–166.
- Zhang, Z., He, Q., Gao, J., Ni, M., 2018. A deep learning approach for detecting traffic accidents from social media data. *Transportation research part C: emerging technologies* 86, 580–596.