# Training (and pre-training) a Neural Network for Structural Model Estimation

Max Wei<sup>1</sup> Zhenling Jiang<sup>2</sup>

<sup>1</sup>Marshall School of Business, University of Southern California

<sup>2</sup>Wharton School, University of Pennsylvania

# Use Machine Learning to Estimate Structural Models

- Three broad types of empirical models:
  - Structural estimation
  - Machine learning
  - Reduced-form analysis
- Today at a high level:
  - Use <u>machine learning</u> techniques to make <u>structural estimation</u> as accessible as reduced form analysis.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

# Neural Networks

- Neural networks are good at predictive tasks.
- Example: image recognition



# Pretrained Model for Image Recognition

• A neural net maps an image to the label "cat".



▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

# This Paper

# Training (and pretraining) a neural network to estimate structural model.

Given your data  $\mapsto$  parameter  $\theta$  (based on a specific model)

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

# Neural Network to Map Dataset to Model Parameter $\theta$

 $\bullet$  For model estimation: treat a dataset as "images", and model parameter  $\theta$  as the label.



▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

# Some Intuitions

- Many structural papers show data patterns (e.g. reduced-form results), which are informative of model parameter  $\theta$ .
  - Higher-ranked items are searched more  $\rightarrow$  search cost increases with ranking.

- ▶ Lower-price items are searched more → utility decreases with price.
- So an "expert" familiar with search model can guesstimate its parameter θ after seeing data patterns.
- We can train such an "**expert**" using neural net.

# Some Intuitions

- Many structural papers show data patterns (e.g. reduced-form results), which are informative of model parameter  $\theta$ .
  - Higher-ranked items are searched more  $\rightarrow$  search cost increases with ranking.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

- Lower-price items are searched more  $\rightarrow$  utility decreases with price.
- So an "expert" familiar with search model can guesstimate its parameter θ after seeing data patterns.
- We can train such an "expert" using neural net.

# Structural Models and Computational Challenges

- Structural model:  $\boldsymbol{y} = \boldsymbol{g}(\boldsymbol{x}, \varepsilon; \boldsymbol{\theta}).$ 
  - E.g., discrete choice, search model.
- $\theta$  is typically estimated by (simulated) MLE/GMM.
  - Often high computational cost because of simulation burden, non-smooth objective functions, etc.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ● ●

Steep learning / coding costs for researchers.

### Overview

- A review of Neural Net Estimator (NNE)<sup>1</sup>
- Extend to a *pre-trained* estimator (**pNNE**)
- Conclusion

<sup>&</sup>lt;sup>1</sup>Yanhao (Max) Wei, Zhenling Jiang (2024) Estimating Parameters of Structural Models Using Neural Networks. Marketing Science 🚊 🔗 🔍

# NNE vs. pNNE

#### pNNE

like using a pre-trained image classifier

- "Off-the-shelf", as easy as running a regression and very fast.
- Trained for a specific structural model.
- Accommodate a wide range of datasets.
- Best for fairly standardized models.

#### NNE

like training own image classifier

- Requires coding by researchers. Compared to SMLE/SMM, can have lighter computational cost and robust to redundant moments.
- Users train for their own (structural) model and specific data.

▲□▶ ▲□▶ ▲□▶ ▲□▶ □ のQで

• A wider range of applications.

# Overview of NNE

• Key idea: use neural nets to learn the mapping  $\{\mathbf{y}_i, \mathbf{x}_i\}_{i=1}^n \mapsto \boldsymbol{\theta}$ .

$$\begin{array}{c} \text{train a} \\ \text{neural net } f(\cdot) \\ \begin{cases} \theta^{(1)} \xrightarrow{\boldsymbol{g}(\boldsymbol{x}_{i},\boldsymbol{e}_{i}^{(1)};\theta^{(1)})} & \{\boldsymbol{y}_{i}^{(1)},\boldsymbol{x}_{i}\}_{i=1}^{n} \xrightarrow{\text{moments}} \boldsymbol{m}^{(1)} \xrightarrow{\text{neural net}} \widehat{\boldsymbol{\theta}}^{(1)} \\ \theta^{(2)} \longrightarrow & \{\boldsymbol{y}_{i}^{(2)},\boldsymbol{x}_{i}\}_{i=1}^{n} \longrightarrow \boldsymbol{m}^{(2)} \longrightarrow \widehat{\boldsymbol{\theta}}^{(2)} \\ & \vdots \\ \theta^{(L)} \longrightarrow & \{\boldsymbol{y}_{i}^{(L)},\boldsymbol{x}_{i}\}_{i=1}^{n} \longrightarrow \boldsymbol{m}^{(L)} \longrightarrow \widehat{\boldsymbol{\theta}}^{(L)} \\ \text{on real data} \\ \end{cases}$$

• Notation:  $\theta^{(\ell)}$  drawn from a space  $\Theta$ ;  $\mathbf{y}_i = \mathbf{g}(\mathbf{x}_i, \varepsilon_i; \theta)$  a structural model;  $\mathbf{y}^{(\ell)}$  simulated outcome;  $\mathbf{m}^{(\ell)}$  simulated moments;  $\hat{\boldsymbol{\theta}}$  neural net prediction.

# Generate Training Datasets

- Draw  $\theta^{(\ell)}$  uniformly from a parameter space  $\Theta$ .
- Use the structural model to generate training datasets.
- Repeat *L* times to get *L* number of training datasets.



・ロト・(理ト・(ヨト・(ヨト・)の())

### Summarize with Data Moments

• Summarize  $\{y_i^{(\ell)}, x_i\}_{i=1}^n$  with data moments  $m^{(\ell)}$ .

- Mean, variance, cross-covariance, etc.
- The parameters need to be identified by the moments (same as in GMM).

$$\begin{array}{c} \text{train a} \\ \text{neural net } f(\cdot) \end{array} \begin{cases} \theta^{(1)} \xrightarrow{\boldsymbol{g}(\boldsymbol{x}_i, \boldsymbol{e}_i^{(1)}; \boldsymbol{\theta}^{(1)})} & \{\boldsymbol{y}_i^{(1)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{\text{moments}} \boldsymbol{m}^{(1)} \xrightarrow{\text{neural net}} \widehat{\boldsymbol{\theta}}^{(1)} \\ \theta^{(2)} \xrightarrow{} & \{\boldsymbol{y}_i^{(2)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{} \boldsymbol{m}^{(2)} \xrightarrow{} \widehat{\boldsymbol{\theta}}^{(2)} \\ & \vdots \\ \theta^{(L)} \xrightarrow{} & \{\boldsymbol{y}_i^{(L)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{} \boldsymbol{m}^{(L)} \xrightarrow{} \widehat{\boldsymbol{\theta}}^{(L)} \end{cases} \end{cases}$$

#### Train a Neural Net $f: \boldsymbol{m} \mapsto \boldsymbol{\theta}$

• Train a neural net to predict  $\theta$  from moments m with MSE loss function  $C = \sum_{\ell=1}^{L} \|\widehat{\theta}^{(\ell)} - \theta^{(\ell)}\|^2.$ 

$$\operatorname{train a}_{neural net f(\cdot)} \begin{cases} \theta^{(1)} \xrightarrow{\boldsymbol{g}(\boldsymbol{x}_i, \boldsymbol{\varepsilon}_i^{(1)}; \theta^{(1)})} & \{\boldsymbol{y}_i^{(1)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{moments} \boldsymbol{m}^{(1)} \xrightarrow{neural net} \widehat{\boldsymbol{\theta}}^{(1)} \\ \theta^{(2)} \xrightarrow{} & \{\boldsymbol{y}_i^{(2)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{} \boldsymbol{m}^{(2)} \xrightarrow{} \widehat{\boldsymbol{\theta}}^{(2)} \\ & \vdots \\ \theta^{(L)} \xrightarrow{} & \{\boldsymbol{y}_i^{(L)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{} \boldsymbol{m}^{(L)} \xrightarrow{} \widehat{\boldsymbol{\theta}}^{(L)} \end{cases}$$

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

#### **Obtain Estimates**

Once the neural net f is trained, plug real data moments to obtain estimates.

 $\begin{array}{c} train \ a \\ neural \ net \ f(\cdot) \end{array} \begin{cases} \theta^{(1)} \xrightarrow{\boldsymbol{g}(\boldsymbol{x}_i, \boldsymbol{e}_i^{(1)}; \theta^{(1)})} & \{\boldsymbol{y}_i^{(1)}, \boldsymbol{x}_i\}_{i=1}^n \xrightarrow{moments} \boldsymbol{m}^{(1)} \xrightarrow{neural \ neural \$ 

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ● ●

# Theoretical Results

• **Proposition**: Under conditions, as the number of training datasets  $L \to \infty$ , the trained neural net converges to limited-info Bayesian posterior mean  $\mathbb{E}(\theta|\mathbf{m})$  and covariance matrix  $\mathbb{C}\text{ov}(\theta|\mathbf{m})$ .

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

- "Limited information" because conditional on *m* not whole data.
- This result holds for any fixed data size n.

# Key Benefits Compared to Existing Methods

- Lighter simulation cost than traditional estimators.
  - No need to evaluate likelihood or moment functions; can be costly (many simulations at every θ trial) and introduce inaccuracy.

- Robust to redundant moments.
  - Learn which moments are useful during training.

# NNE Recap

• Train a neural net to estimate structural models.

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

- Code available at nnehome.github.io.
- User still needs to
  - 1) simulate  $\boldsymbol{y}_i = \boldsymbol{g}(\boldsymbol{x}_i, \varepsilon_i; \boldsymbol{\theta})$
  - 2) generate moments
  - 3) train a neural net

#### Overview

- A review of Neural Net Estimator (NNE)
- Extend to a pre-trained estimator (pNNE)

(ロ)、(型)、(E)、(E)、(E)、(O)へ(C)

Conclusion

# Pretrained Estimator

- Benefit: Requires *minimal* coding and computation time of users.
- Scope:
  - ► Trained for a specific structural model (e.g., consumer search).
  - Applies to a wide range of data.
- Key innovations on top of NNE:
  - ▶ In training, **x** is simulated instead of taken from data.
  - ▶ *m* includes reduced-form regressions in addition to summary moments.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00

• Build a pretrained estimator for sequential search model as a proof of concept.

#### Consumer Sequential Search Model

• Consumer optimal sequential search model (Weitzman 1979).

- Consumer utility:  $u_{ij} = \beta' \mathbf{x}_{ij}^{(prod)} + e_{ij} + \epsilon_{ij}$
- Outside option:  $u_{i0} = \eta_0 + \eta' \mathbf{x}_i^{(cons)} + \epsilon_{i0}$
- Consumer pay search cost  $c_{ij}$  to know  $\epsilon_{ij}$ :  $log(c_{ij}) = \alpha_0 + \alpha' \mathbf{x}_{ij}^{(ads)}$
- Optimal search strategy: which options to search, when to stop, and what to buy.

- Observe consumer search and purchase decisions.
- Estimate parameters:  $\{\mathbf{y}, \mathbf{x}\} \mapsto [\alpha_0, \boldsymbol{\alpha}, \eta_0, \boldsymbol{\eta}, \boldsymbol{\beta}]$ .

# Estimation via MLE

• Very challenging to estimate: millions of search-buy combinations for each consumer:

$$\begin{split} L &= \prod_{i=1}^{N} \int \left\{ z_{i1} > z_{i2} > \ldots > z_{iH} > \max_{j' \notin (1, \ldots, H)} (z_{ij'}), \\ \max(u_{i0}, u_{i1}) < z_{i2}, \ldots, \ \max(u_{i0}, u_{i1}, \ldots, u_{i, H-1}) < z_{iH}, \ \max(u_{i0}, u_{i1}, \ldots, u_{iH}) > \max(z_{ij'}), \\ j' \notin (1, \ldots, H) \\ u_{ij^{\star}} &\geq \max(u_{i0}, u_{i1}, \ldots, u_{iH}) \right\} \int_{j=1}^{J} dF(\epsilon_{ij}) \int_{j=0}^{J} dF(\epsilon_{ij}) dF(\epsilon_{ij}) \\ \end{split}$$

- Directly simulating likelihood is nearly impossible.
- Survey by Ursu, Seiler and Honka 2024<sup>2</sup>

<sup>&</sup>lt;sup>2</sup>Ursu, Raluca and Seiler, Stephan and Honka, Elisabeth (2024), The Sequential Search Model: A Framework for Empirical Research. Quantitative Marketing and Economics.

# pNNE Demo

- Off-the-shelf: use on your data.
- (Beta-version) available at pnnehome.github.io.
- Python package coming very soon!



▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

# Monte Carlo

- Benchmark: MLE with a GHK-type estimator (Jiang et al. 2021<sup>3</sup>).
- pNNE has lower RMSE than MLE with the number of simulations S = 1000.
- pNNE is much faster (4-5 orders of magnitude difference).



<sup>&</sup>lt;sup>3</sup> Jiang, Zhenling, Tat Chan, Hai Che, and Youwei Wang (2021). "Consumer search and purchase: An empirical investigation of retargeting based on consumer online behaviors." Marketing Science.

# Monte Carlo: Fourteen Datasets from Different Domains

	number of	product	consumer	"advertising"
	consumers	attributes	attributes	attributes
laobao				
iPad mini	10,000	3	5	0
iPad 4	10,000	3	5	0
Expedia (Kaggle)				
Destination 1	1,258	6	2	1
Destination 2	8,97	6	2	1
Expedia (WCA)				
Cancun	2,135	6	0	1
Manhattan	1,369	6	0	1
Trivago				
PC channel	2,082	7	0	1
mobile channel	2,726	7	0	1
JD				
PC channel	5,896	4	2	0
mobile channel	4,069	4	2	0
app channel	10,000	4	2	0
Anonymous eCommerce				
laptop	10,000	6	0	0
vacuum	10,000	4	0	0
washer	10,000	5	0	0

# Monte Carlo Results

■: MLE (R=50) ■: MLE (R=1e3) ■: pretrained NNE



# Real Data with Diverse Search Patterns

	number of	buy rate	search rate	number of
	consumers			searcnes
Taobao				
iPad mini	10,000	1.16%	10.82%	1.14 1.12
iPad 4	10,000	1.23%	9.58%	
Expedia (Kaggle)				
Destination 1	1,258	11.53%	8.66%	1.18
Destination 2	8,97	3.79%	6.02%	1.13
Expedia (WCA)				
Cancun	2,135	1.31%	30.82%	1.56
Manhattan	1,369	3.07%	28.71%	1.52
Trivago				
PC channel	2,082	2.83%	40.01%	2.03
mobile channel	2,726	2.46%	29.09%	1.62
JD				
PC channel	5,896	16.35%	26.46%	1.44
mobile channel	4,069	11.16%	12.17%	1.15
app channel	10,000	7.96%	12.88%	1.17
Anonymous eCommerce				
laptop	10,000	5.9%	21.5%	1.36
vacuum	10,000	8.5%	19.0%	1.29
washer	10,000	7.8%	22.1%	1.38

#### Real Data Estimation Results

■: MLE (R=1e3)

: pretrained NNE



◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 - のへで

# Overview of pNNE

• Generate training examples, construct data patterns, and train a neural net to predict  $\theta$ .

$$\begin{array}{c} train \\ pNNE \\ \begin{cases} \theta^{(1)}, \mathbf{x}^{(1)} \xrightarrow{\mathbf{g}(\mathbf{x}_{i}^{(1)}, \mathbf{e}_{i}^{(1)}; \theta^{(1)})} & \{\mathbf{y}_{i}^{(1)}, \mathbf{x}_{i}^{(1)}\}_{i=1}^{n} \xrightarrow{reduced form} \mathbf{m}^{(1)} \xrightarrow{neural net} \widehat{\theta}^{(1)} \\ \\ \theta^{(2)}, \mathbf{x}^{(2)} \xrightarrow{} & \{\mathbf{y}_{i}^{(2)}, \mathbf{x}_{i}^{(2)}\}_{i=1}^{n} \xrightarrow{} \mathbf{m}^{(2)} \xrightarrow{} \widehat{\theta}^{(2)} \\ & \vdots \\ \\ \theta^{(L)}, \mathbf{x}^{(L)} \xrightarrow{} & \{\mathbf{y}_{i}^{(L)}, \mathbf{x}_{i}^{(L)}\}_{i=1}^{n} \xrightarrow{} \mathbf{m}^{(L)} \xrightarrow{} \widehat{\theta}^{(L)} \\ \end{array}$$

$$apply to \\ real data \\ \begin{cases} \mathbf{y}_{i}, \mathbf{x}_{i} \\ real data \end{cases}^{n} \xrightarrow{reduced form} \mathbf{m} \xrightarrow{neural net} \underbrace{\widehat{\theta}}_{estimate} \\ \end{cases}$$

#### Overview of pNNE: Generate x

Generate diverse x to cover possible values in real applications.



# How to Generate Diverse x?

• Draw the data dimensionality from a range.

Specified Ranges for Data Dimensions							
	d <sup>(prod)</sup>	d <sup>(ads)</sup>	d <sup>(cons)</sup>	J	п		
min	2	0	0	15	1e3		
max	8	2	5	35	$\infty$		

▲□▶ ▲□▶ ▲□▶ ▲□▶ □ のQで

- Draw from correlated multivariate normal distributions.
- Randomly transform to different non-normal distributions.
  - Dummy
  - A scale of 1 to 5
  - A skewed continuous distribution
- Standardize each attribute.

# Overview of pNNE: Draw $\theta$

- Draw  $\theta$  from a prior distribution.
  - Working with standardized x is very helpful.
  - ▶ Need the prior to be wide enough to cover possible values but not too wide.



# Overview of pNNE: Generate Training Data

- With  $\mathbf{x}^{(l)}$  and  $\boldsymbol{\theta}^{(l)}$ , use the structural model to generate  $\mathbf{y}^{(l)}$ .
- Repeat *L* times to get *L* number of training datasets.



・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・

#### Overview of pNNE: Construct Data Patterns

Summarize the key information with *m*.



・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・
 ・

# How to Construct Informative Data Patterns?

- Regression coefficients:
  - Search decisions: Logit regression of y<sup>(search)</sup> on product and consumer attributes.
  - Purchase decisions: Multinomial logit regression of y<sup>(buy)</sup><sub>ij</sub> among searched products.

- Similar regressions at the consumer level.
- Summary statistics:
  - Buy rate, search rate, average number of searches.

### Overview of pNNE: Train a Neural Net

• Train a neural net to predict  $\theta$  from moments m.

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

### Overview of pNNE: Obtain Estimates

• Users plug in their data to the pretrained model.





# "Privacy-preserving" Structural Estimation

 Researchers only need to know the aggregated data moments *m* to estimate the model.



company compiles moments

# Conclusion

- One can train (and use a pre-trained) neural network model to estimate structural models.
- Key idea: learn the mapping from reduced-form patterns to parameters.
- pNNE is trained for a specific structural model (consumer search as a proof of concept) while NNE can be applied to a wide set of applications.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

# Potential Applications of pNNE

- Privacy sensitive scenarios: Estimation without access to individual-level data.
  - The (aggregated) reduced-form patterns serve as sufficient statistics for the structural model.
- A fast "package" for model estimation.
  - Scalable to larger data.
  - Larger models where search model is a component and needs to be estimated repeated (e.g., dynamic models).
  - Incorporate search model into algorithms (e.g., Bandit) that are updated in real time.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00