


AI-guided Surrogate Modeling for Real-world Optimization

Yuandong Tian
Research Scientist and Manager

Meta AI (FAIR)



Reinforcement Learning



Go



Chess



Shogi



Poker



DoTA 2



StarCraft II


Big Success in Games

ELF OpenGo

Vs top professional players


Name (rank)	ELO (world rank)	Result
Kim Ji-seok	3590 (#3)	5-0
Shin Jin-seo	3570 (#5)	5-0
Park Yeonghun	3481 (#23)	5-0
Choi Cheolhan	3466 (#30)	5-0

Single GPU, 80k rollouts, 50 seconds
Offer unlimited thinking time for the players



What's Beyond Games?

Chip Design (Google)




Several weeks with **human experts** in the loop




Fully automatic design in 6 hours


Optimization Problems




Travel Salesman Problem




Job Scheduling



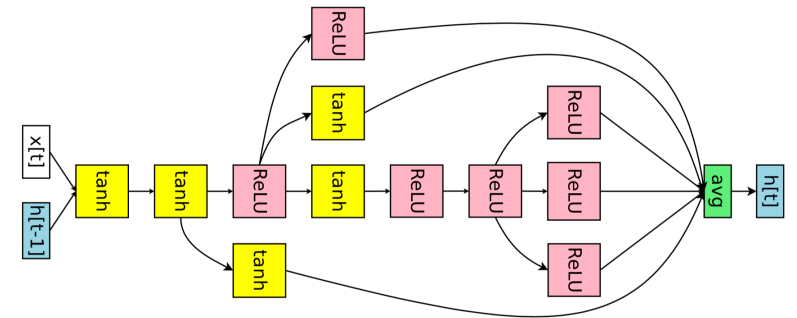
Vehicle Routing



Bin Packing



Protein Folding



Model-Search


$$x^* = \arg \max_{x \in \Omega} f(x)$$

Wait...What?



- Many problems are NP-hard problems.
 - No good algorithm unless $P = NP$
- These guarantees are worst-case ones.
 - To prove a lower-bound, construct an adversarial example to fail the algorithm
- For specific distribution, there might be better heuristics.
 - Human heuristics are good but may not be suitable for everything

Efficient Search for Games

Go




Chess



~~Human Knowledge~~


Machine learned models

Efficient Search for Optimization



Exhaustive search to get **a good solution**

Human Knowledge



More Efficient Search for Optimization

Can we use
Machine Learning?

Exhaustive search to get **a good solution**


Components of Search

Search Heuristics

State Representation

Design of State/Action Space

Design Surrogate Models



Components of Search



Search Heuristics

Initial solution → Improved solution1 →
Improved solution2 → ...

[X. Chen and Y. Tian, **Learning to Perform Local Rewriting for Combinatorial Optimization**, *NeurIPS'19*]

[H. Shi et al, **Deep Symbolic Superoptimization Without Human Knowledge**, *ICLR'20*]

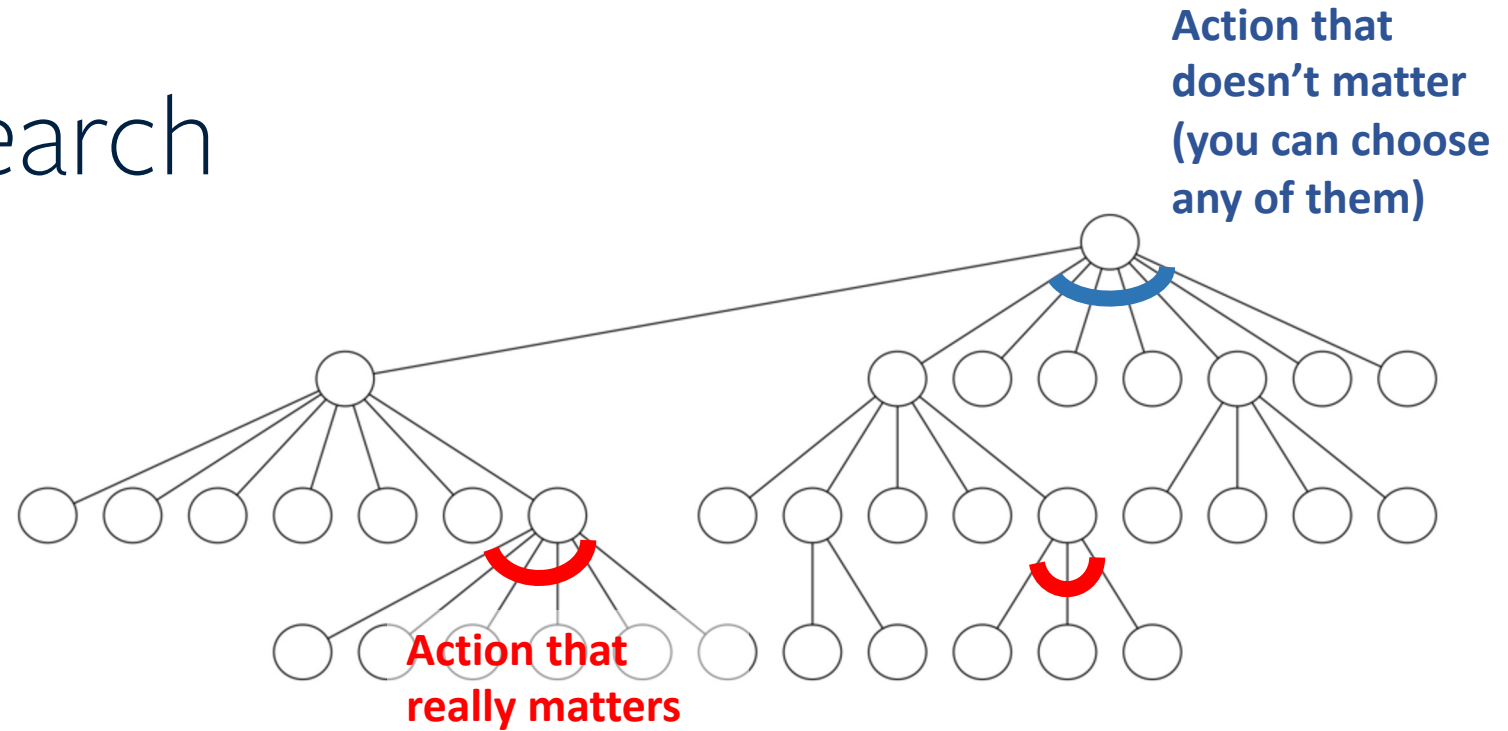
[H. Zhu et al, **Network planning with deep reinforcement learning**, *SIGCOMM'21*]

[T. Huang et al, **Local Branching Relaxation Heuristics for Integer Linear Programs**, *CPAIOR'23*]

[T. Huang et al, **Searching Large Neighborhoods for Integer Linear Programs with Contrastive Learning**, *arXiv'23*]

Components of Search

Design of
State/Action Space



If useful actions only happen after **50** binary moves, then we will waste our efforts in this 2^{50} possibilities.

[L. Wang, et al, **Learning Search Space Partition for Black-box Optimization using MCTS**, *NeurIPS'20*]

[L. Wang, et al, **Sample-Efficient Neural Architecture Search by Learning Action Space**, *T-PAMI'21*]


[K. Yang, et al, **Learning Space Partitions for Path Planning**, *NeurIPS'21*]

[Y. Zhao, et al, **Multi-objective Optimization by Learning Space Partitions**, *ICLR'22*]

[Y. Liang, et al, **Learning Compiler Pass Orders using Coreset and Normalized Value Prediction**, *arXiv'23*]

Components of Search

State
Representation



GOAL: Letting in as much sunlight as possible

[X. Chen et al, **Latent Execution for Neural Program Synthesis Beyond Domain-Specific Languages**, *NeurIPS'21*]

[T. Wang et al, **Denoised MDPs: Learning World Models Better Than the World Itself**, *ICML'22*]

[Z. Jiang et al, **Efficient Planning in a Compact Latent Action Space**, *ICLR'23*]

Components of Search

Design Surrogate
Models

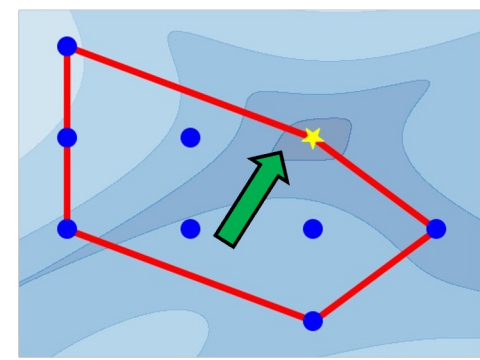
*[Y. Zhao, et al, **Few-shot neural architecture search**, ICML'21]*

*[D. Zha, et al, **DreamShard: Generalizable Embedding Table Placement for Recommender Systems**, NeurIPS'22]*

*[A. Ferber, et al, **SurCo: Learning Linear Surrogates For Combinatorial Nonlinear Optimization Problems**, arXiv'22]*

*[A. Cohen, et al, **Modeling Scattering Coefficients using Self-Attentive Complex Polynomials with Image-based Representation**, arXiv'23]*

SurCo: Learning Linear Surrogates for Combinatorial Nonlinear Optimization



Aaron Ferber¹, Taoan Huang¹, Daochen Zha², Martin Schubert³, Benoit Steiner⁴, Bistra Dilkina¹, Yuandong Tian⁴

¹University of Southern California, ²Rice University, ³Reality Lab Display, ⁴Meta AI (FAIR)




RICE UNIVERSITY
School of Engineering
Department of Computer Science




Optimizing Nonlinear Functions over Combinatorial Regions

- Nonlinear + differentiable objective
- Combinatorial feasible region
- Real-world domains:
 - Computer system planning
 - Designing photonic devices
 - Throughput optimization
 - Antenna design
 - Energy grid




Example: Embedding Table Placement




Given:

- k tables
- n identical devices
- Table i has memory requirement m_i
- Device j has memory capacity M_j

Naive Sharding




Balanced Sharding



Find


- Allocation of tables to devices observing device memory limits
- Minimize latency which is **estimated by a neural network** (capturing nonlinear interactions)

Example: Embedding Table Placement



Given:

- k tables
- n identical devices
- Table i has memory requirement m_i
- Device j has memory capacity M_j

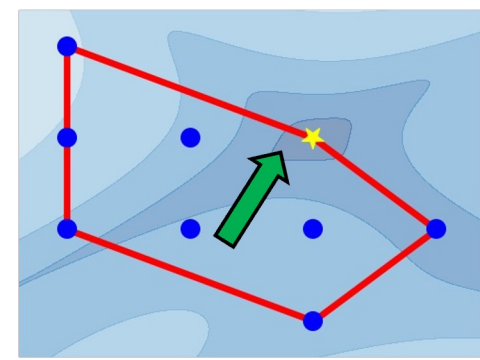


Formulation

$$\text{Min}_x \mathbf{L}(\{x_{ij}\}) \quad \text{s.t.} \quad \sum_i x_{ij} m_i \leq M_j, \quad \sum_j x_{ij} = 1, \quad x_{ij} \in \{0,1\}$$

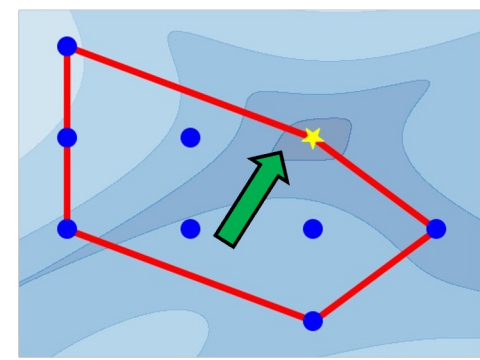
\mathbf{L} is nonlinear due to system issues (e.g., batching, communication, etc)

Nonlinear Optimization is **Hard**



- Specific domains have specialized solvers
- General solvers are often slow (without very careful modeling)
- Genetic algorithms or gradient-based methods may not find feasible solutions

Linear Optimization is **Easy**(ish)



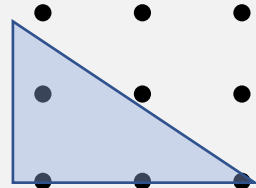
- MILP solvers (CPLEX, Gurobi, SCIP) easily handle industry-scale problems
- Plus other solvers for linear settings
 - Greedy
 - LP + total unimodularity

Idea: Find a Linear Surrogate


- Learn a MILP objective whose optimal solution x^* solves the nonlinear problem

Originally

Nonlinear optimization with combinatorial constraints

$$\begin{aligned} \min_x f(x; y) \\ \text{s.t. } x \in \Omega = \end{aligned}$$


combinatorial constraints



Now

Surrogate optimization

$$\begin{aligned} x^*(y) = \operatorname{argmin}_x c(y)^T x \\ \text{s.t. } x \in \Omega \end{aligned}$$


solved by existing combinatorial solvers

Predict surrogate cost $c = c(y)$



Idea: Find a Linear Surrogate


- Learn a MILP objective whose optimal solution x^* solves the nonlinear problem



Challenge: how to find the right objective?

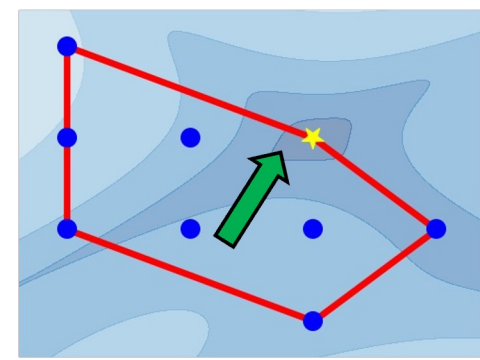
Idea: Find a Linear Surrogate

- Learn a MILP objective whose optimal solution x^* solves the nonlinear problem




Proposal: gradient-based optimization


Proposal: surrogate learning



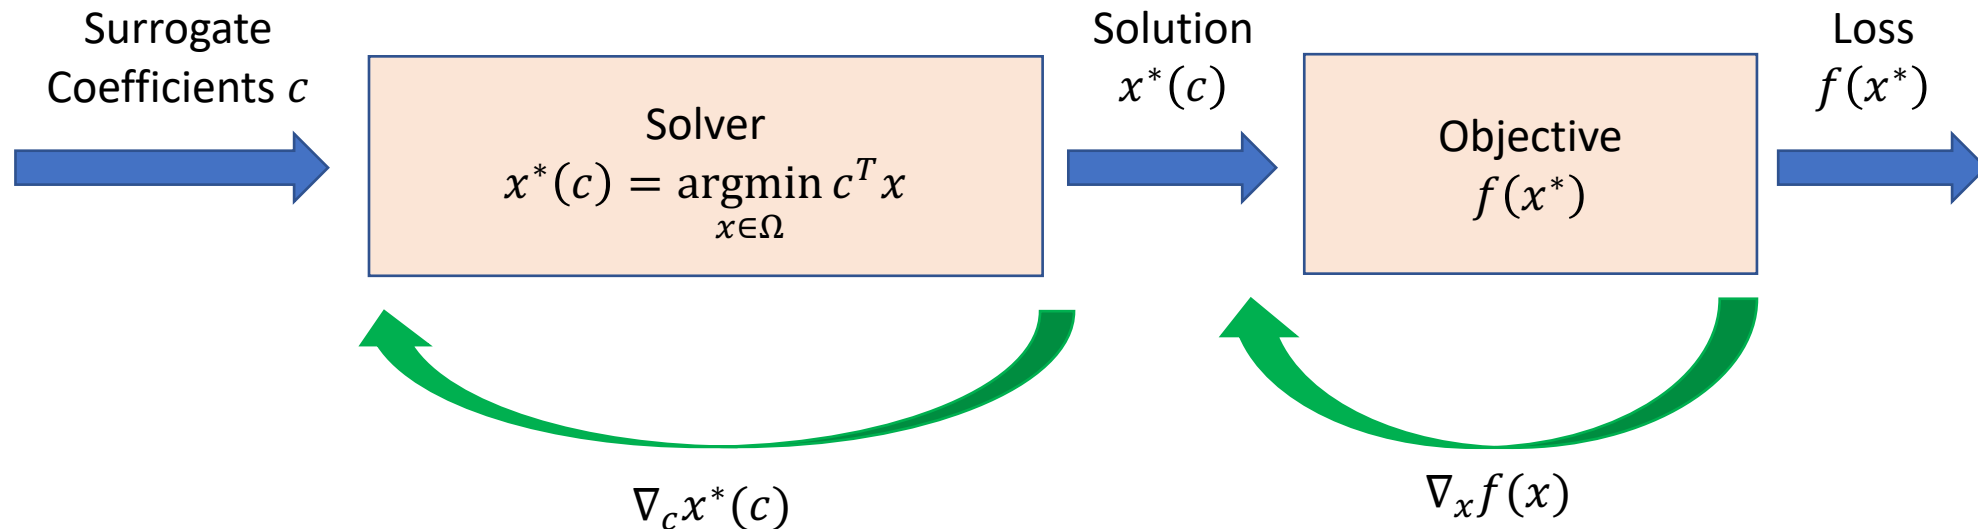
- Use surrogate MILP to solve original problem
- Find linear coefficients c such that $\operatorname{argmin}_{x \in \Omega} f(x) \approx \operatorname{argmin}_{x \in \Omega} c^T x$



SurCo-zero: gradient-based optimization



- **Iterative** solver based on linear surrogate guided by **gradient updates**
- Update linear coefficients c such that $x^*(c)$ improves objective $f(x^*(c))$



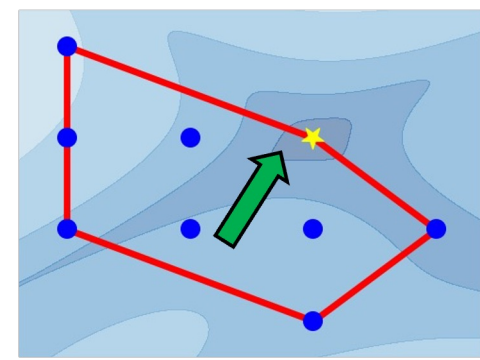
Recent work on differentiable optimization

Differentiation of blackbox optimizers

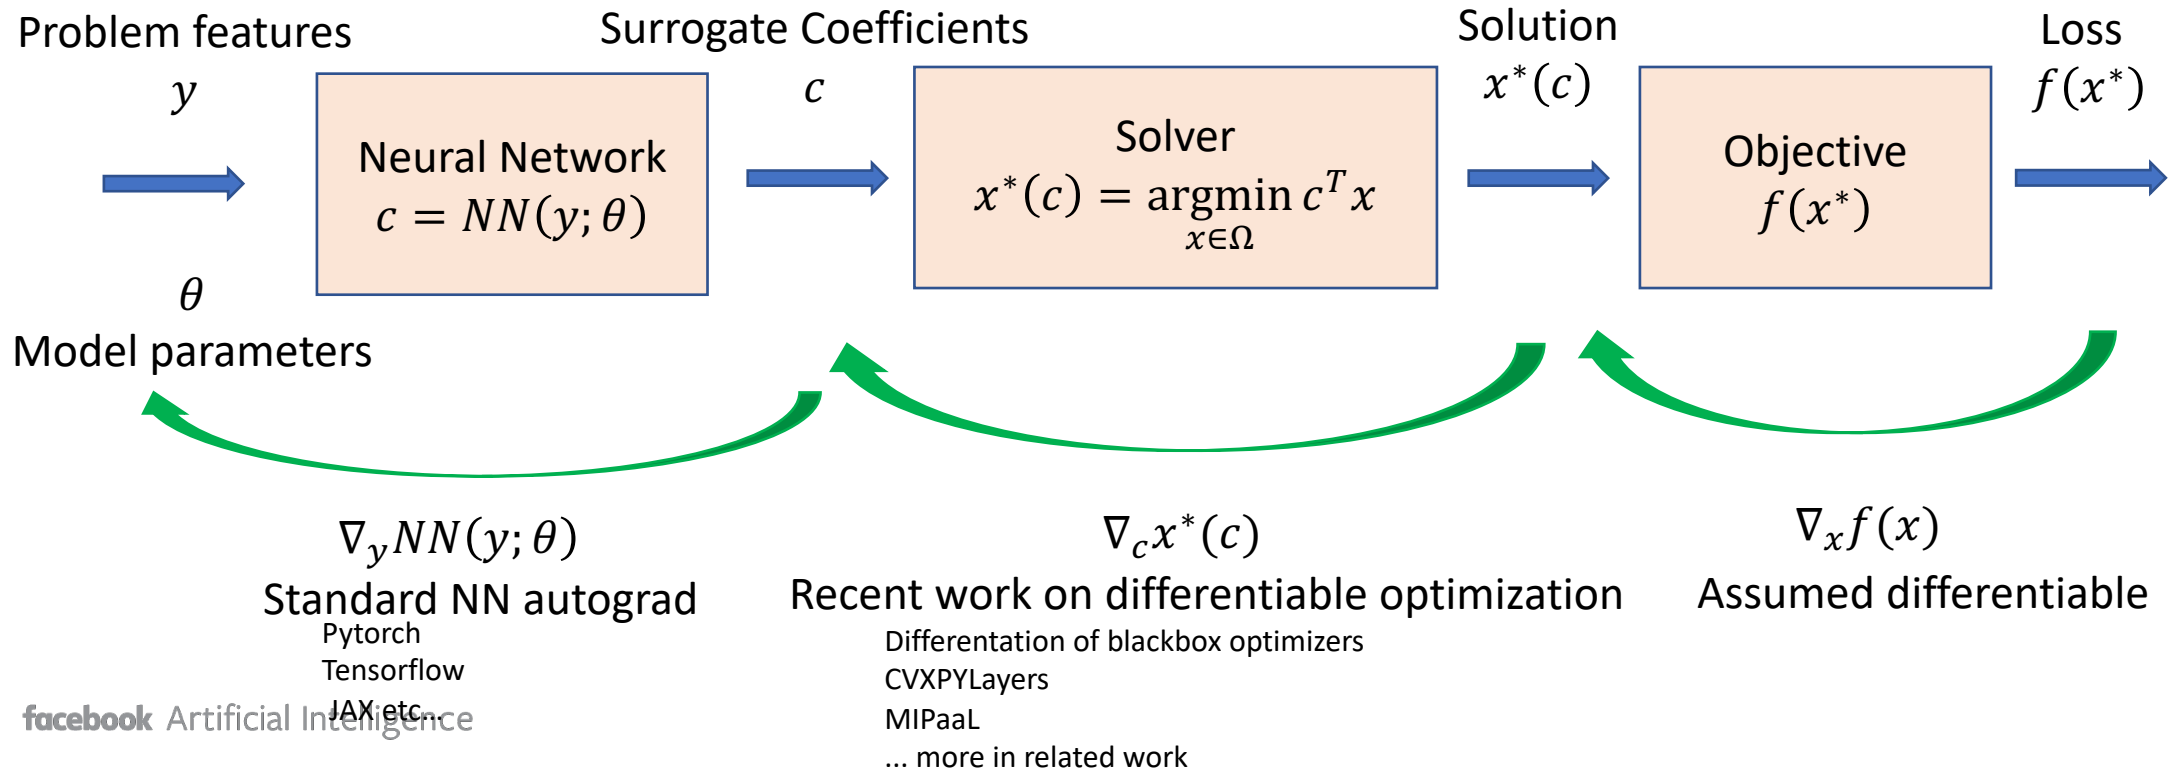
CVXPYLayers

MIPaal


SurCo-prior: distributional learning



- One pass solver based on model **learned offline**
- Use neural model based on **problem features** to predict linear coefficients




SurCo-prior: distributional learning



- Update neural network parameters from training dataset


$$c_i = NN(y_i; \theta)$$



Train Model parameters θ

Surrogate Coefficients

$$c_{\text{test}} = NN(y_{\text{test}}; \theta)$$




Solution

$$x^*(c_{\text{test}})$$




SurCo-hybrid: fine-tuning from trained model



Update neural network parameters from training dataset

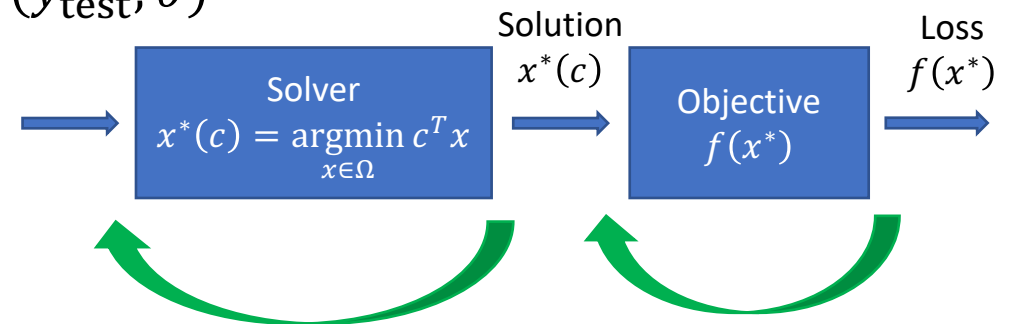
Fine-tune surrogate on-the-fly

$$c_i = NN(y_i; \theta)$$



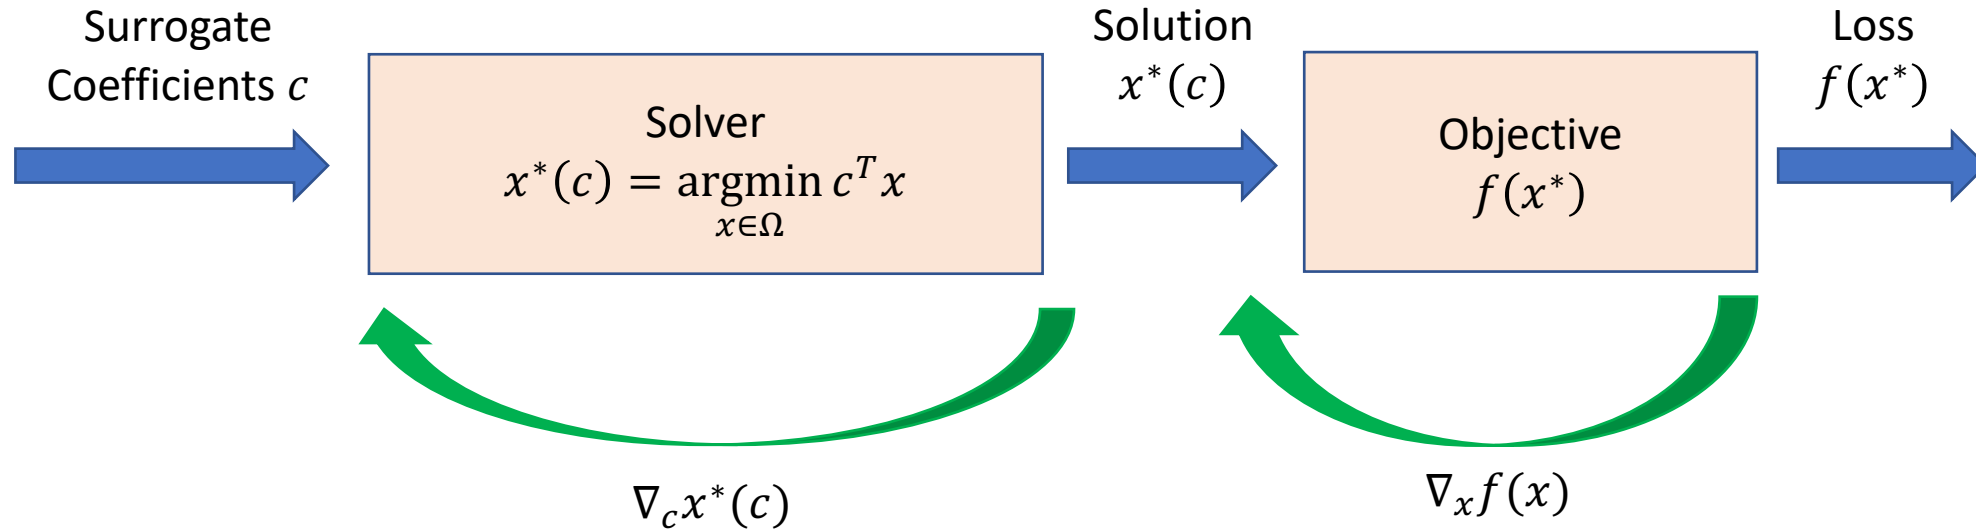

Initial Surrogate Coefficients

$$c_0 = NN(y_{\text{test}}; \theta)$$



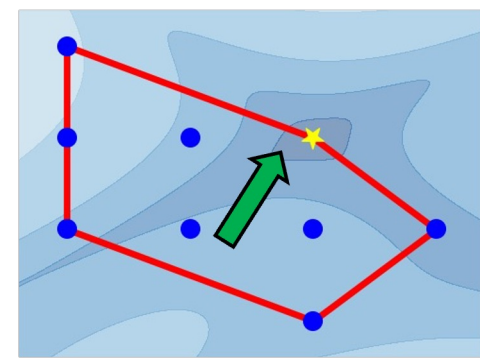
Train Model parameters θ

SurCo-zero




No offline training data, just solve a single problem instance on-the-fly

SurCo-prior




$$c_i = NN(y_i; \theta)$$



Train Model parameters θ

Surrogate Coefficients

$$c_{\text{test}} = NN(y_{\text{test}}; \theta)$$




Solution

$$x^*(c_{\text{test}})$$






Uses offline training data to quickly solve problems at test time with just one solver call

SurCo-hybrid



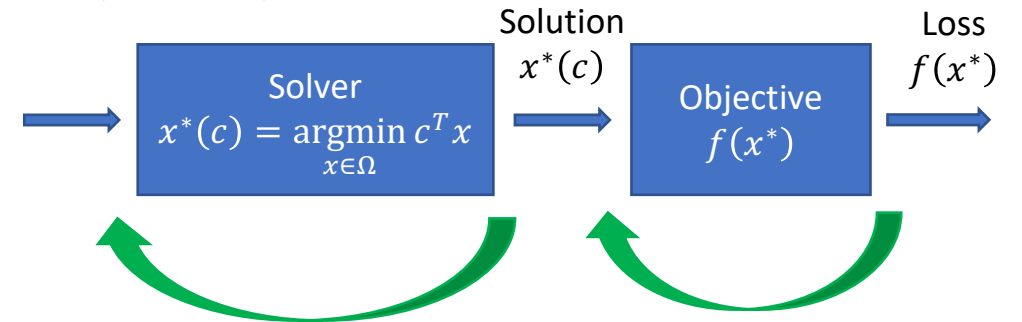
$$c_i = NN(y_i; \theta)$$



Train Model parameters θ

Initial Surrogate Coefficients

$$c_0 = NN(y_{\text{test}}; \theta)$$



Offline train + on-the-fly fine-tuning the surrogate

Related Work

Differentiable optimization: backprop through solvers

Amos et al. OptNet: Differentiable optimization as a layer in neural networks. ICML 2017

Agrawal et al. Differentiable Convex Optimization Layers. NeurIPS 2019

Berthet et al. Learning with Differentiable Perturbed Optimizers. NeurIPS 2020

Demirović et al. Predict+Optimise with Ranking Objectives: Exhaustively Learning Linear Functions. IJCAI 2019

Demirović et al. Dynamic Programming for Predict + Optimise. AAI 2020

Djulonga et al. Differentiable Learning of Submodular Models. NeurIPS 2017

Donti et al. Task-Based End-to-End Model Learning in Stochastic Optimization. NeurIPS 2017

Elmachtoub et al. Smart “Predict, then Optimize”. Management Science 2022

Ferber et al. MIPaaL: Mixed Integer Program as a Layer. AAI 2020

Lee et al. Meta-Learning with Differentiable Convex Optimization. CVPR 2019

Mandi et al. Smart Predict-and-Optimize for Hard Combinatorial Optimization Problems. AAI 2020

Niepert et al. Implicit MLE: Backpropagating Through Discrete Exponential Family Distributions. NeurIPS 2021

Valstelica et al. Differentiation of Blackbox Combinatorial Solvers. ICLR 2019

Rolnšek et al. Optimizing Rank-Based Metrics with Blackbox Differentiation. CVPR 2020

Wang et al. Automatically Learning Compact Quality-Aware Surrogates for Optimization Problems. NeurIPS 2020


Wang et al. SATNet: Bridging Deep Learning and Logical Reasoning Using a Differentiable Satisfiability Solver. ICML 2019

Wilder et al. Melding the Data-Decisions Pipeline: Decision-focused Learning for Combinatorial Optimization. AAI 2019

Wilder et al. End to End Learning and Optimization on Graphs. NeurIPS 2019


How SurCo is different from Predict+Optimise?

Predict+Optimize



- Suitable for linear optimization problems.
- Requires a ground truth linear coefficients c_i of the objective

SurCo



- Suitable for **nonlinear** objective $f(x; y)$.
- Unlike existing nonlinear solvers, **NO** analytic form needed.
- Does **NOT** require a ground truth linear coefficients c_i . Learned surrogate coefficients by itself.

Both requires contextual information (i.e., problem description y_i)

Related Work

Mixed Integer Nonlinear Optimization: general-purpose solvers

Burer et al. Non-Convex Mixed Integer Nonlinear Programming: A Survey. ORMS 2012

Belotti et al. Mixed Integer Nonlinear Optimization. Acta Numerica 2013

General-purpose heuristic optimizers: combinatorial constraints are hard

Gad et al. Pygad: An Intuitive Genetic Algorithm Python Library. 2021

Rapin et al. Nevergrad – A Gradient-Free Optimization Platform. 2018

Wang et al. Learning Search Space Partition for Black-Box Optimization Using Monte Carlo Tree Search. NeurIPS 2020

Wang et al. Sample Efficient Neural Architecture Search by Learning Actions for Monte Carlo Tree Search. PAMI 2021

RL for combinatorial optimization: combinatorial constraints are hard

Khalil et al. Learning Combinatorial Optimization Algorithms Over Graphs. NeurIPS 2017

Kool et al. Attention, Learn to Solve Routing Problems! ICLR 2018

Mazyavkina et al. Reinforcement Learning for Combinatorial Optimization: A Survey. COR 2021

Nazari et al. Reinforcement Learning for Solving the Vehicle Routing Problem. NeurIPS 2018

Zhang et al. A Reinforcement Learning Approach to Job-Shop Scheduling. IJCAI 1995

Embedding Table Sharding

Used in large-scale deep learning systems: recommendation systems, knowledge graph

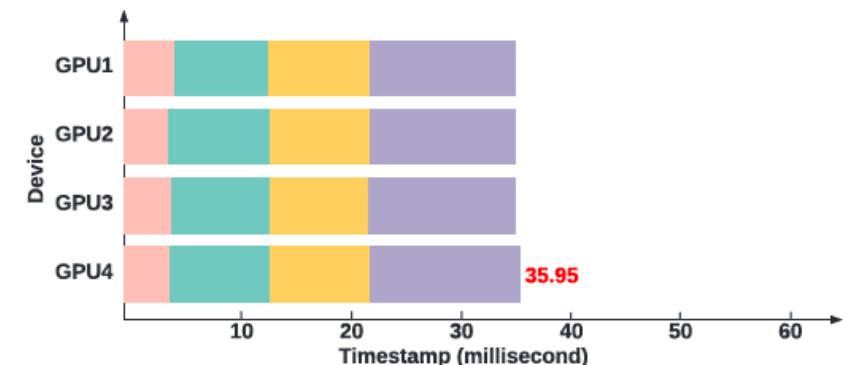

Place N “tables” (with known memory need m_i) on K devices ($x_{ij} = 1$: table i assigned to device j)

$$\text{Min}_x L(\{x_{ij}\}) \quad \text{s.t.} \quad \sum_i x_{ij} m_i \leq M_j, \quad \sum_j x_{ij} = 1, \quad x_{ij} \in \{0,1\}$$


L : Runtime bottleneck $f(x)$ estimated by NN (longest-running device)

L is nonlinear due to system issues
(e.g., batching, communication, etc.)

$c(y; \theta)$ gives surrogate “per-table cost” c_{ij}
(and $\sum_{ij} c_{ij} x_{ij}$ is the surrogate latency objective)



Embedding Table Sharding





- Public **D**eep **L**earning **R**ecommendation **M**odel (DLRM dataset) placing between 10 to 60 tables on 4 GPUs
- Baseline: Greedy
- SoTA: RL approach Dreamshard¹
- SurCo: Surrogate NN model learned via CVXPYLayers (differentiable LP Solver)




¹ Zha et al. NeurIPS 2022

Dataset: https://github.com/facebookresearch/dlrm_datasets

Results – Table Sharding




Inverse Photonic Design



- Design physically-viable devices that take light waves and routes different wavelengths to correct locations

$$\mathcal{L}(S) = \left(\left\| \text{softplus} \left(g \frac{|S|^2 - |S_{\text{cutoff}}|^2}{\min(w_{\text{valid}})} \right) \right\|_2 \right)^2$$

- Device design misspecification loss $f(x)$ computed by differentiable electromagnetic simulator
- Feasible solution: the design must be the union of brush pattern
 - $x = \text{binary_opening}(x, \text{brush})$
 - $x = \sim \text{binary_opening}(\sim x, \text{brush})$



Inverse Photonic Design


- Dataset: Ceviche Challenges¹
- Most baselines don't work here due to combinatorial constraints
- SoTA: Brush-based algorithm¹

- SurCo: Surrogate learned via blackbox differentiation² of brush solver


¹Schubert et al. ACS Photonics 2022

²Vlastelica et al. ICLR 2019


Dataset: <https://github.com/google/ceviche-challenges>




Waveguide bend



Beam splitter





Mode converter





Wavelength division multiplexer


Results – Inverse Photonics




Inverse photonics Convergence comparison + Solution example



E_z magnitude
first wavelength



E_z magnitude
second wavelength




Wavelength division multiplexer

Takeaways:

- SurCo-Zero finds loss-0 solutions quickly
- SurCo-Hybrid uses offline training data to get a head start

Conclusion



- Handle industrial applications with differentiable optimization
- High-quality solutions to combinatorial nonlinear optimization by finding linear surrogates
 - Sometimes we can find “easier” surrogate problems that solve much more difficult instances
- SurCo works in several data settings
 - Zero-shot vs Offline training
 - One step inference vs fine-tuning

Sample-efficient Surrogate Model for Frequency Response of Linear PDEs using Self-Attentive Complex Polynomials

Andrew Cohen^{1*}, Weiping Dou², Jiang Zhu², Slawomir Koziel³, Peter Renner²,
Jan-Ove Mattson², Xiaomeng Yang¹, Beidi Chen^{1,4}, Kevin Stone¹, Yuandong Tian^{1*}

¹Meta AI (FAIR), ²Reality Lab Antenna (Meta), ³Reykjavik University, ⁴Carnegie Mellon University

* = Equal technical contribution




Background

Solving the linear PDE

$$\frac{\partial^n \psi}{\partial t^n} = F(\psi, \nabla_x \psi, \dots; \mathbf{h})$$

$\psi(x, t)$ is the spatial-temporal signal under time evolutions.
 F is a linear function with respect to ψ and its derivatives
 \mathbf{h} is design choice.




Examples




$$\frac{\partial \psi}{\partial t} = \nabla^2 \psi$$

Heat equation




$$i\hbar \frac{\partial \psi}{\partial t} = \left[-\frac{\hbar^2}{2m} \nabla^2 + V \right] \psi$$

Schrodinger's Equation



$$\frac{\partial^2 \psi}{\partial t^2} = c^2 \nabla^2 \psi$$

Wave equation



$$\begin{aligned} \nabla \cdot E &= \frac{\rho}{\epsilon_0}, \nabla \cdot B = 0 \\ \nabla \times E &= -\frac{\partial B}{\partial t}, \nabla \times B = \mu_0 j + \frac{1}{c^2} \frac{\partial E}{\partial t} \end{aligned}$$

Maxwell's equation


Tricky to simulate accurately and efficiently → Can we do better?

Antenna Design problem


Goal:

find the right design to achieve the right frequency response

$h =$



$S_{11}(\psi) =$




Antenna Design problem

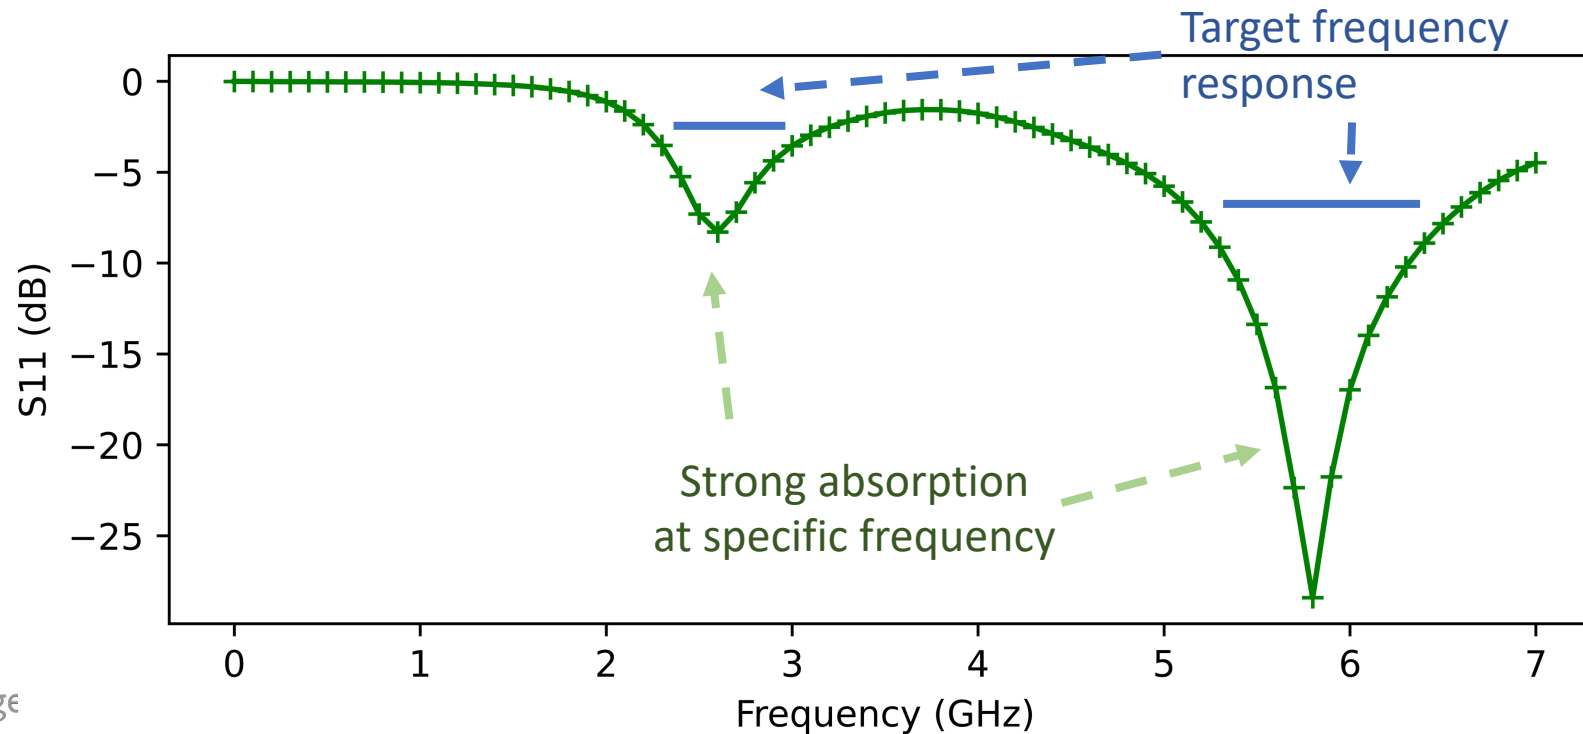
Goal:

find the right design to achieve the right frequency response

$h =$



$S_{11}(\psi) =$



Discretization of linear PDE systems

$$\frac{\partial^n \psi}{\partial t^n} = F(\psi, \nabla_x \psi, \dots; \mathbf{h}) \quad \longrightarrow \quad \frac{\partial \phi}{\partial t} = A(\mathbf{h})\phi$$

The matrix A encodes the information of F

One example: $\frac{\partial^2 \psi}{\partial t^2} = c^2 \nabla^2 \psi \quad \longrightarrow \quad \phi = \left[\psi(x_1), \dots, \psi(x_N), \frac{\partial \psi}{\partial t}(x_1), \dots, \frac{\partial \psi}{\partial t}(x_N) \right]$

Discretized onto N vertices

$$A = \begin{bmatrix} 0 & I \\ c^2 B & 0 \end{bmatrix}$$


Frequency Domain

Signal in the temporal domain: $\phi(x, t)$

Vector form $\boldsymbol{\phi}(t) = [\phi(x_1, t), \dots, \phi(x_N, t)]$

Signal in the frequency domain: $\hat{\phi}(x, \omega) = \int \phi(x, t) e^{-i\omega t} dt$

Vector form $\hat{\boldsymbol{\phi}}(\omega) = [\hat{\phi}(x_1, \omega), \dots, \hat{\phi}(x_N, \omega)]$



Parametric formula for Linear PDEs

Theorem: For any linear coefficients \mathbf{b}_1 and \mathbf{b}_2 :

$$\frac{\mathbf{b}_1^T \widehat{\boldsymbol{\phi}}(\omega)}{\mathbf{b}_2^T \widehat{\boldsymbol{\phi}}(\omega)} = c_0(\mathbf{h}) \prod_{k=1}^{K_1} (\omega - \mathbf{z}_k(\mathbf{h})) \prod_{k=1}^{K_2} (\omega - \mathbf{p}_k(\mathbf{h}))^{-1}$$

where the constant $c_0(\mathbf{h})$, zeros $\mathbf{z}_k(\mathbf{h})$ and poles $\mathbf{p}_k(\mathbf{h})$ are complex functions of the design choice \mathbf{h}

Proof idea: Linear ODE theory gives us the analytic form of the solution $\boldsymbol{\phi}(t) = e^{At} \boldsymbol{\phi}(0)$. Fourier Transform yields $\widehat{\boldsymbol{\phi}}(\omega)$ as a rational function of complex polynomials w.r.t. frequency ω .

For Antenna Optimization

The *Scattering Coefficients* $S_{11}(\omega)$:

$$S_{11}(\omega) = \frac{Z_{\text{in}}(\omega) - Z_0}{Z_{\text{in}}(\omega) + Z_0}$$

$Z_{\text{in}}(\omega)$: Input *Impedance*. Impedance $Z(\omega) := V(\omega)/I(\omega)$



Both are linear function w.r.t. signal $\hat{\phi}(\omega)$

Parameter form of $\log|S_{11}(\omega)|$

$$\log|S_{11}(\omega)| = \log|c_0(\mathbf{h})| + \sum_{k=1}^K \log \frac{|\omega - \mathbf{z}_k(\mathbf{h})|}{|\omega - \mathbf{p}_k(\mathbf{h})|}$$

where the constant $c_0(\mathbf{h})$, zeros $\mathbf{z}_k(\mathbf{h})$ and poles $\mathbf{p}_k(\mathbf{h})$ are complex functions of the design choice \mathbf{h} .

CZP model architecture





Image-based Representation
of design choice \mathbf{h}



Spatial Attention
(Wu et al. 2020)
+
Transformer
(Vaswani et al. 2017)

$c_0(\mathbf{h})$



$\{z_k(\mathbf{h})\}$

$\{p_k(\mathbf{h})\}$

Predict the constant, zeros and poles from an
image representation of an antenna

Data Collection

Dataset is collected from commercial simulators (e.g., CST)
Numerical simulation of Electromagnetic wave dynamics




It takes minutes (or even hours) to get one simulation data point.
Dataset size = 48k


Model Evaluation

- Static Evaluation
 - On a held-out test set, compute the loss
 - Loss = gap between surrogate models and ground truth (commercial software)
- Dynamic Evaluation
 - Use the surrogate model to search a good design
 - Evaluate the design in ground truth (commercial software)


Static Evaluation: Surrogate Model Test Loss



Static Evaluation: Surrogate Model Test Loss




Static Evaluation: Surrogate Model Test Loss




Visualization

Our CZP model captures the smooth structure of scattering coefficients $S_{11}(\omega)$



(a) CZP



(b) Raw


Dynamic Evaluation: CZP model with Search

We use Soft Actor-Critic as the specific search technique.

Goal: to find a solution to satisfy the frequency constraints (verified with CST)

3 models trained x 3 search attempts using different random seed

For each attempt, check top-3 solutions



Future Work

- The formulation applies to general linear PDEs
 - Maxwell's Equations
 - Schrodinger's Equations
 - Many more ...
- Test on more complicated scenarios.
 - 3D antenna

Thanks!