

# Building Adaptive Networks with Machine Learning

---

**Tushar Swamy**

Stanford University


**Networks need data-driven  
decisions:**

**————→ *Machine Learning***

# Networks can benefit from machine learning



**Anomaly Detection**  
**Traffic Classification**



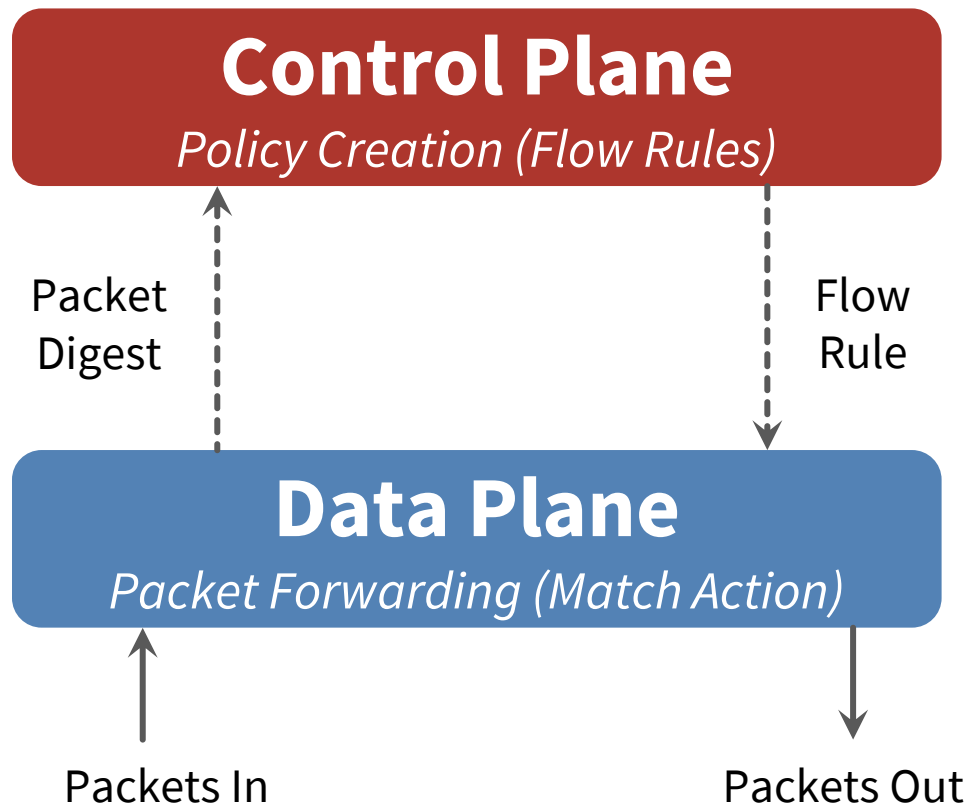
**Congestion Control**  
**Active Queue Management**



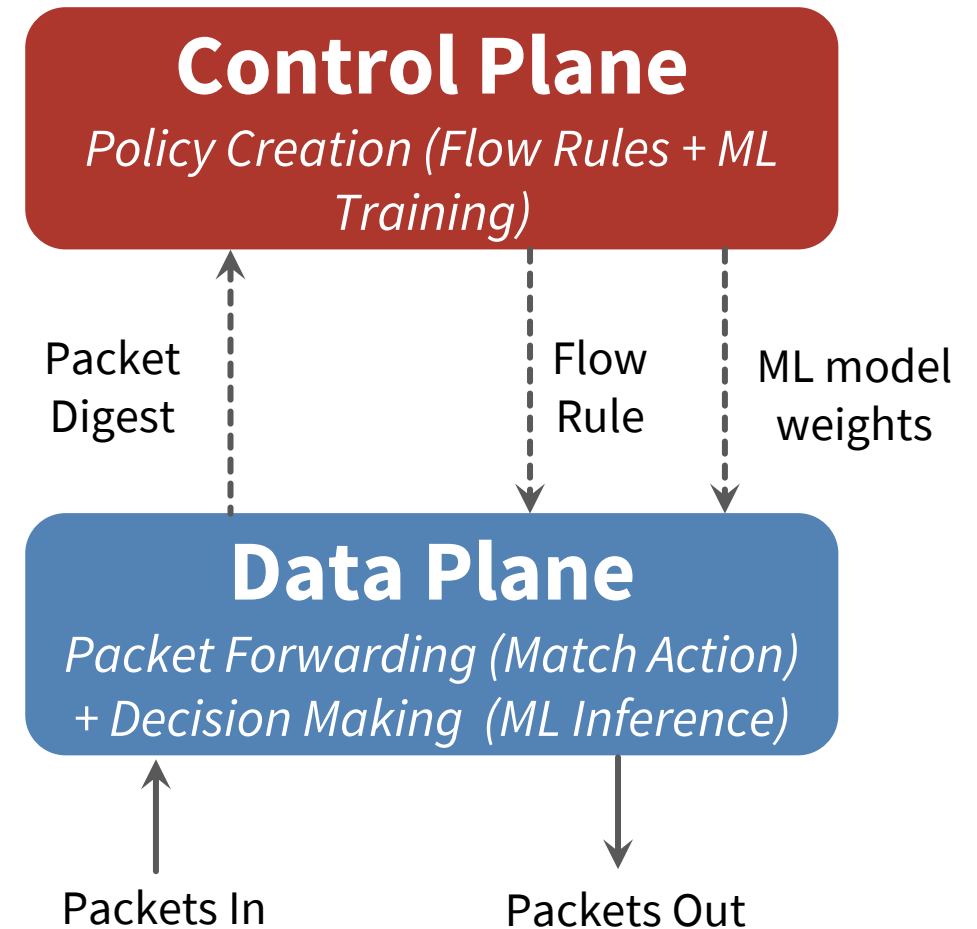
**Queue Reconstruction**  
**QoS Correlation**

# A Taurus network introduces ML for management

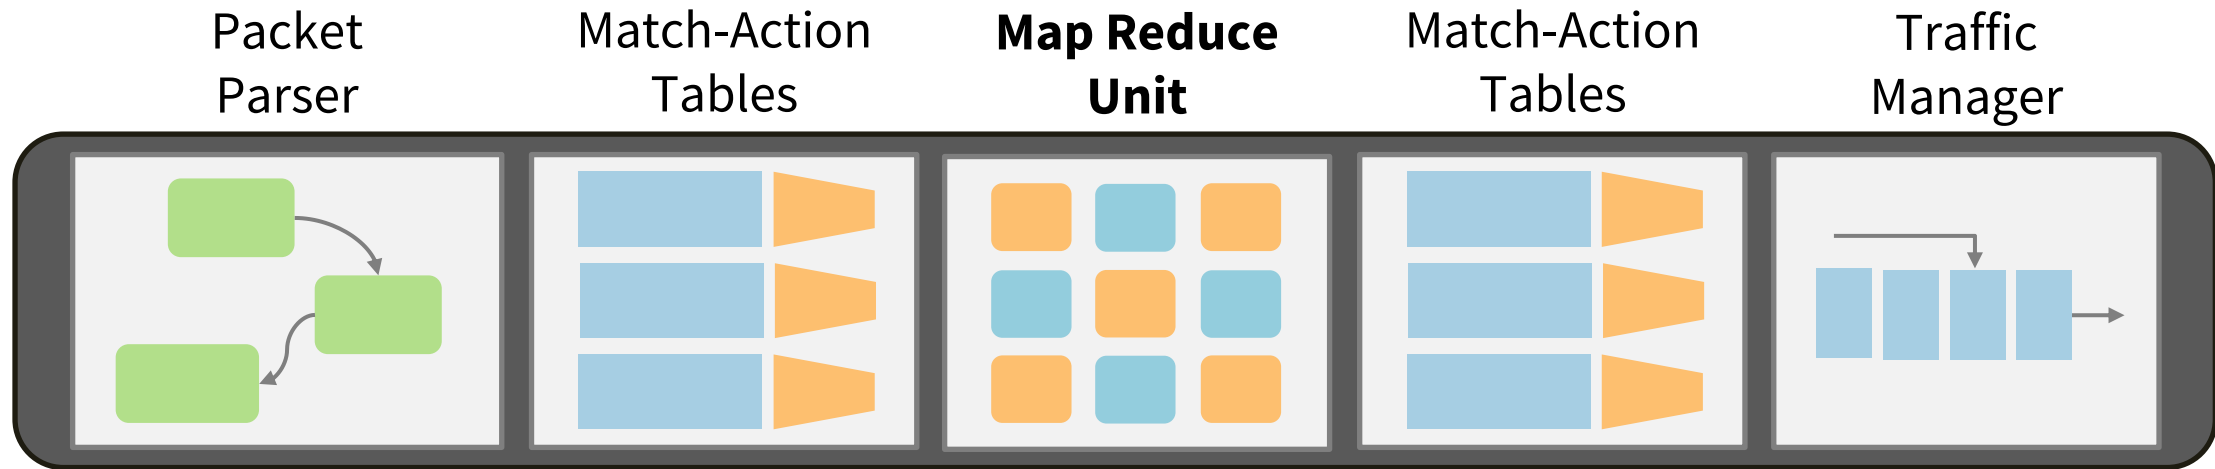
## *Software Defined Network*



## *Software Defined Network with **Taurus***



# The Taurus switch pipeline




# Taurus achieves line-rate ML inference

- Robustness and performance of the network is determined by *quality and speed of reaction*
- ML inference should happen *per-packet* in the *dataplane*
- *Taurus* enables line-rate, per-packet ML inference in the dataplane

*Taurus: a data plane architecture for per-packet ML*  
(ASPLOS '22) <https://dl.acm.org/doi/10.1145/3503222.3507726>

***How do we  
program Taurus?***

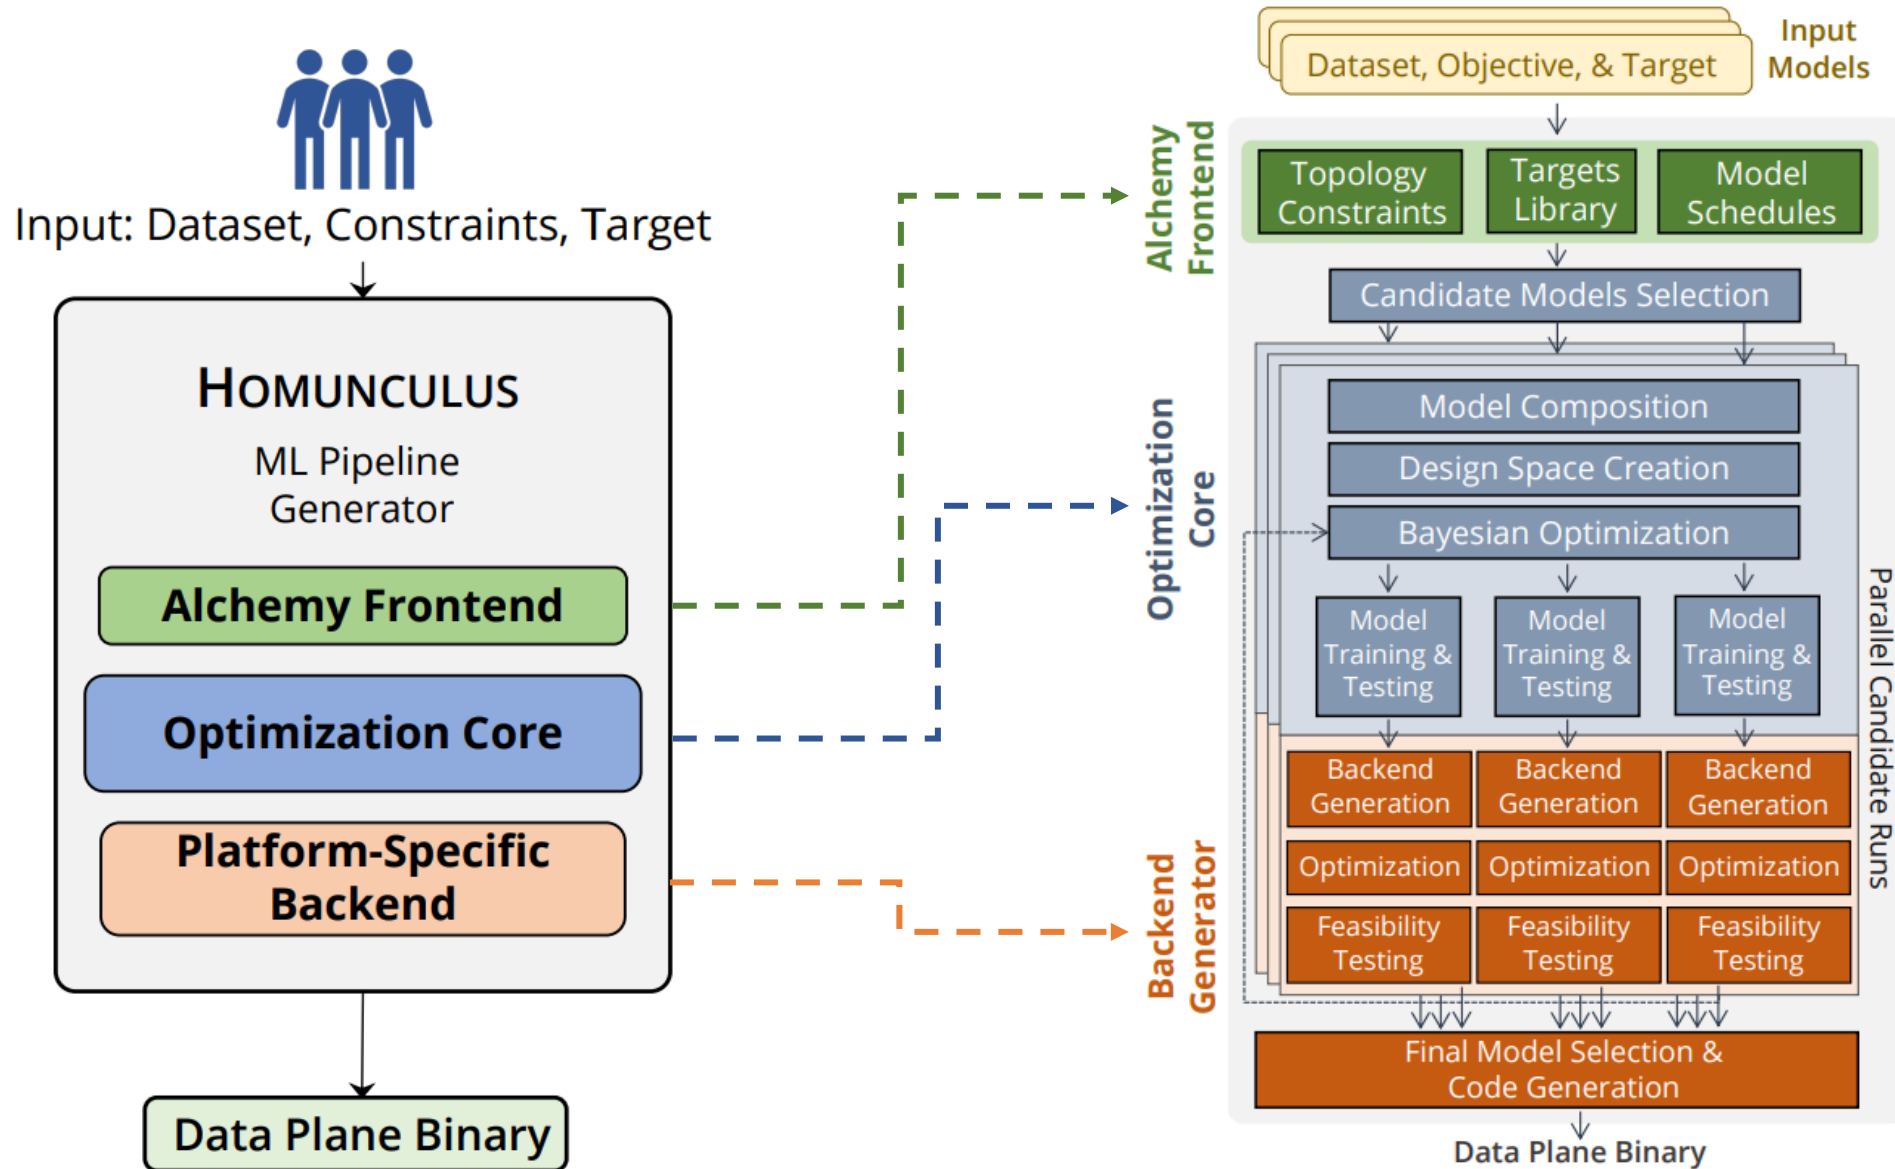


# Homunculus: a framework for dataplane model generation

- Provide **high-level directives** to express user intent
- Take **network and resource constraints** into account when building ML models
- **Generate binaries** for different dataplane devices with optimized ML models

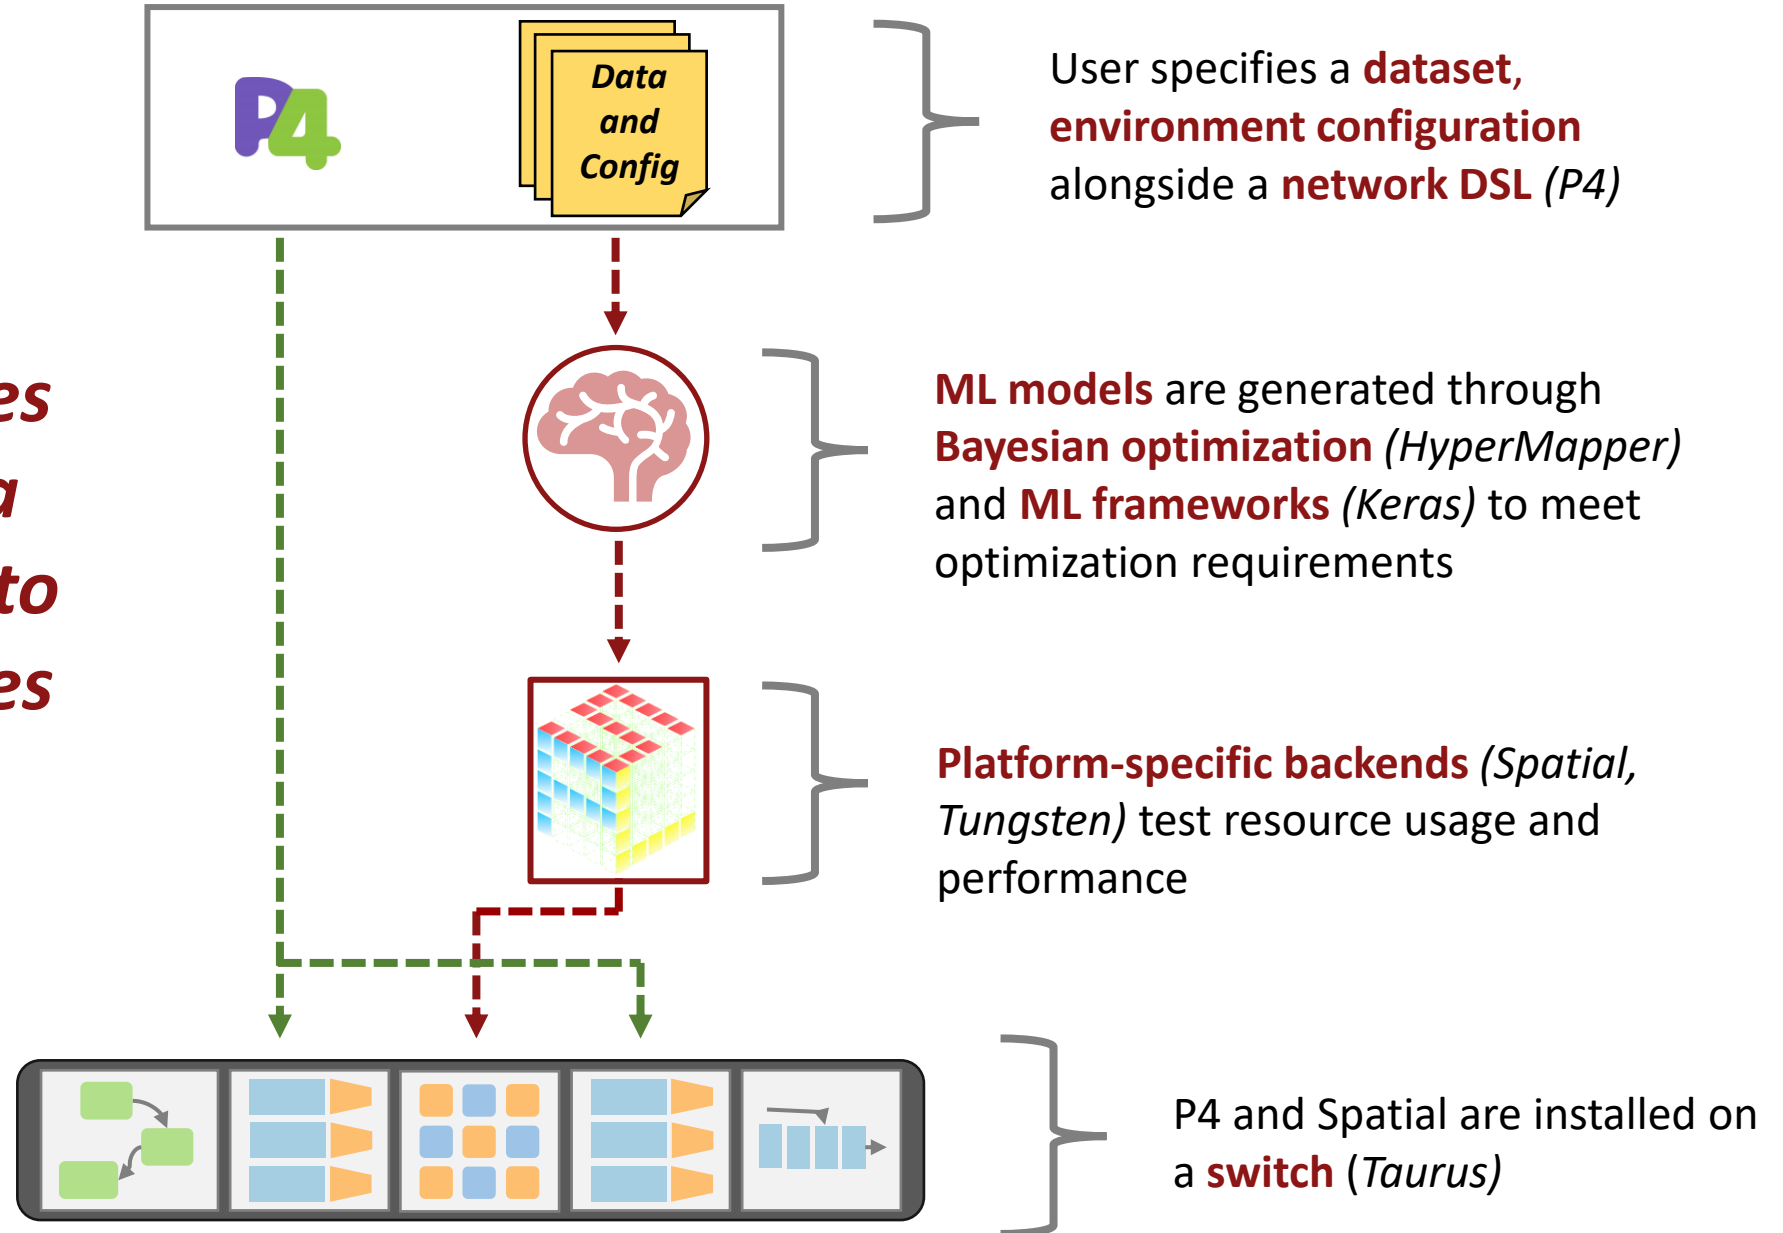


# Homunculus architecture

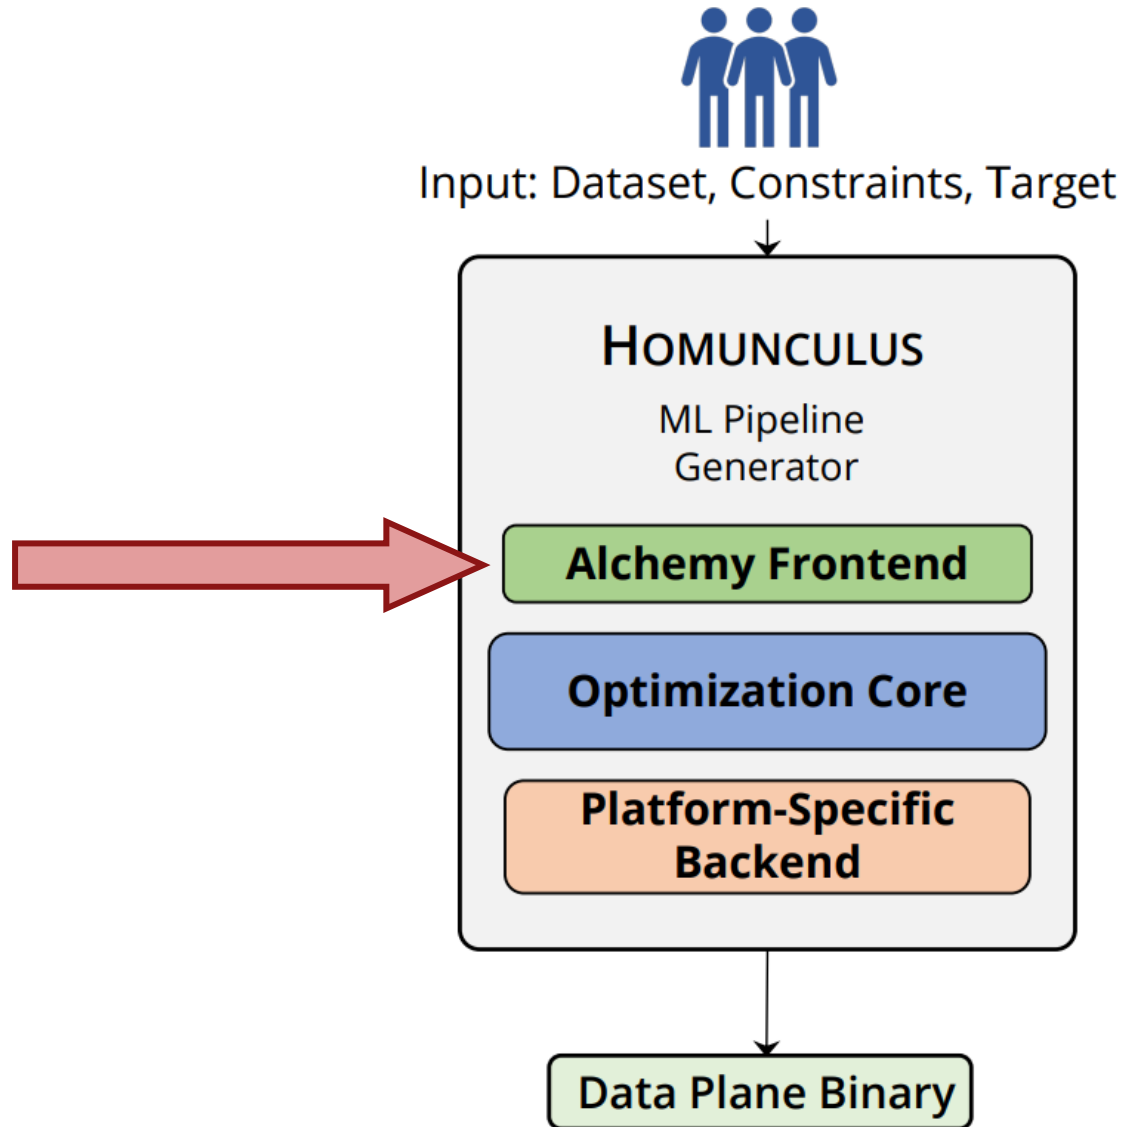


# Homunculus overview

*Homunculus provides network operators a high-level compiler to program ML switches*



# Homunculus frontend: Alchemy



# Homunculus frontend: Alchemy

```
1 import homunculus
2 from homunculus.alchemy import
3     DataLoader, Model, Platforms
4 import ad_loader
5
6 @DataLoader # training data loader definition
7 def wrapper_func():
8     tnx, tny = ad_loader.load_from_file(
9         "train_ad.csv")
10    tsx, tsy = ad_loader.load_from_file(
11        "test_ad.csv")
12    return {
13        "data": {"train": tnx, "test": tsx },
14        "labels": {"train": tny, "test": tsy }}
15
16 # Specify the model of choice
17 model_spec = Model({
18     "optimization_metric": ["f1"],
19     "algorithm": ["dnn"],
20     "name": "anomaly_detection",
21     "data_loader": wrapper_func })
22
23 # Load platform
24 platform = Platforms.Taurus()
25 platform.constrain(
26     "performance": {
27         "throughput": 1, # GPkt/s
28         "latency": 500 }, # ns
29     "resources": { "rows": 16, "cols": 16 })
30
31 # Schedule model and generate code
32 platform.schedule(model_spec)
33 homunculus.generate(platform)
```

← *Alchemy is a Python library*

# Homunculus frontend: Alchemy

```
1 import homunculus
2 from homunculus.alchemy import
3     DataLoader, Model, Platforms
4 import ad_loader
5
6 @DataLoader # training data loader definition
7 def wrapper_func():
8     tnx, tny = ad_loader.load_from_file(
9         "train_ad.csv")
10    tsx, tsy = ad_loader.load_from_file(
11        "test_ad.csv")
12    return {
13        "data": {"train": tnx, "test": tsx },
14        "labels": {"train": tny, "test": tsy }}
15
16 # Specify the model of choice
17 model_spec = Model({
18     "optimization_metric": ["f1"],
19     "algorithm": ["dnn"],
20     "name": "anomaly_detection",
21     "data_loader": wrapper_func })
22
23 # Load platform
24 platform = Platforms.Taurus()
25 platform.constrain(
26     "performance": {
27         "throughput": 1, # GPkt/s
28         "latency": 500 }, # ns
29     "resources": { "rows": 16, "cols": 16 })
30
31 # Schedule model and generate code
32 platform.schedule(model_spec)
33 homunculus.generate(platform)
```



*Load data*

# Homunculus frontend: Alchemy

```
1 import homunculus
2 from homunculus.alchemy import
3     DataLoader, Model, Platforms
4 import ad_loader
5
6 @DataLoader # training data loader definition
7 def wrapper_func():
8     tnx, tny = ad_loader.load_from_file(
9         "train_ad.csv")
10    tsx, tsy = ad_loader.load_from_file(
11        "test_ad.csv")
12    return {
13        "data": {"train": tnx, "test": tsx },
14        "labels": {"train": tny, "test": tsy }}
15
16 # Specify the model of choice
17 model_spec = Model({
18     "optimization_metric": ["f1"],
19     "algorithm": ["dnn"],
20     "name": "anomaly_detection",
21     "data_loader": wrapper_func })
22
23 # Load platform
24 platform = Platforms.Taurus()
25 platform.constrain(
26     "performance": {
27         "throughput": 1, # GPkt/s
28         "latency": 500 }, # ns
29     "resources": { "rows": 16, "cols": 16 })
30
31 # Schedule model and generate code
32 platform.schedule(model_spec)
33 homunculus.generate(platform)
```

← *Specify a model*

# Homunculus frontend: Alchemy

```
1 import homunculus
2 from homunculus.alchemy import
3     DataLoader, Model, Platforms
4 import ad_loader
5
6 @DataLoader # training data loader definition
7 def wrapper_func():
8     tnx, tny = ad_loader.load_from_file(
9         "train_ad.csv")
10    tsx, tsy = ad_loader.load_from_file(
11        "test_ad.csv")
12    return {
13        "data": {"train": tnx, "test": tsx },
14        "labels": {"train": tny, "test": tsy }}
15
16 # Specify the model of choice
17 model_spec = Model({
18     "optimization_metric": ["f1"],
19     "algorithm": ["dnn"],
20     "name": "anomaly_detection",
21     "data_loader": wrapper_func })
22
23 # Load platform
24 platform = Platforms.Taurus()
25 platform.constrain(
26     "performance": {
27         "throughput": 1, # GPkt/s
28         "latency": 500 }, # ns
29     "resources": { "rows": 16, "cols": 16 })
30
31 # Schedule model and generate code
32 platform.schedule(model_spec)
33 homunculus.generate(platform)
```

 *Constrain the platform*

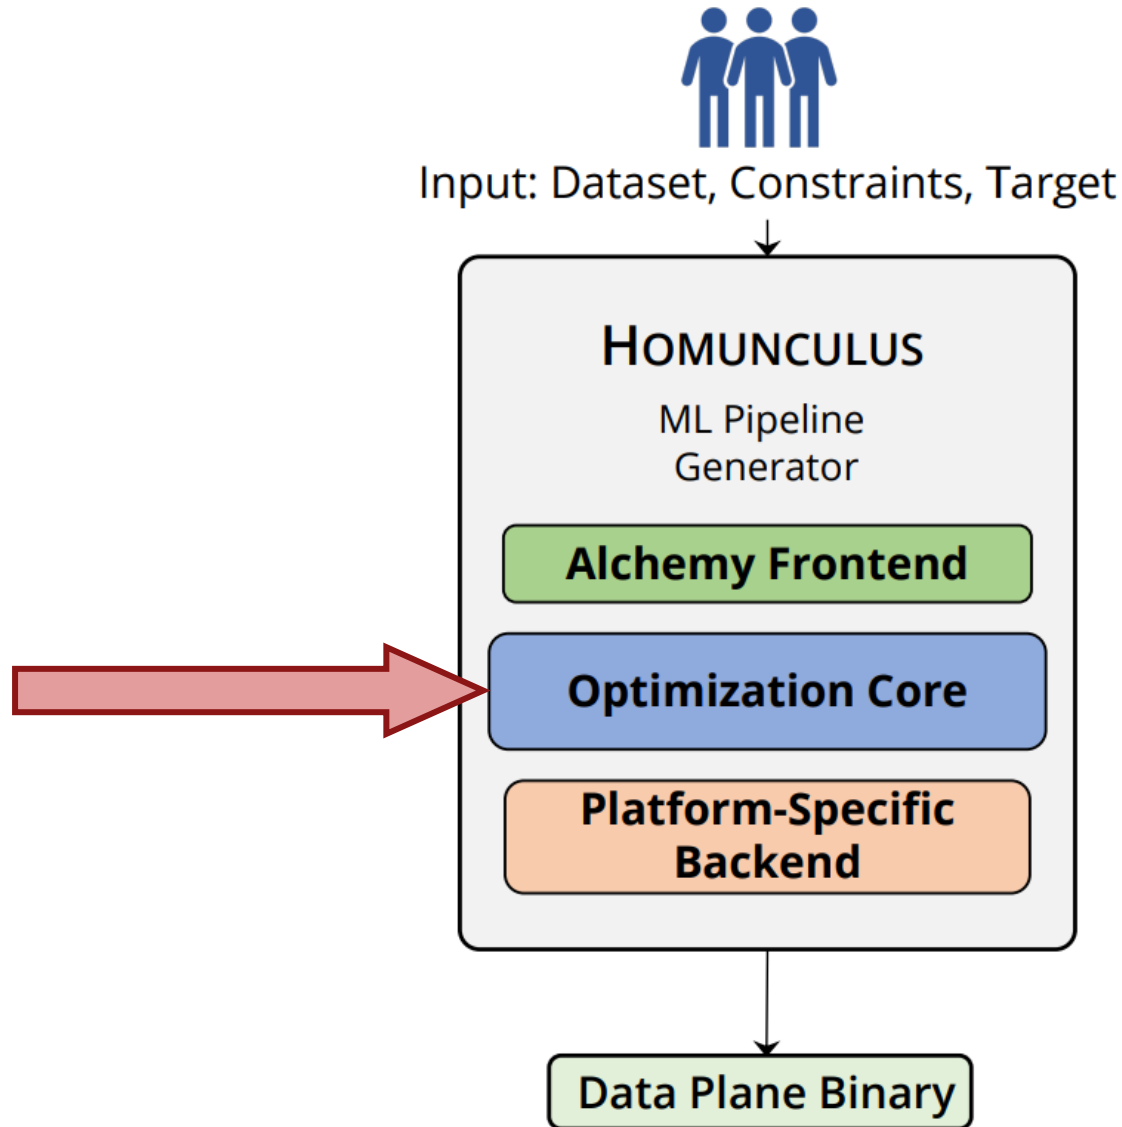
# Homunculus frontend: Alchemy

```
1 import homunculus
2 from homunculus.alchemy import
3     DataLoader, Model, Platforms
4 import ad_loader
5
6 @DataLoader # training data loader definition
7 def wrapper_func():
8     tnx, tny = ad_loader.load_from_file(
9         "train_ad.csv")
10    tsx, tsy = ad_loader.load_from_file(
11        "test_ad.csv")
12    return {
13        "data": {"train": tnx, "test": tsx },
14        "labels": {"train": tny, "test": tsy }}
15
16 # Specify the model of choice
17 model_spec = Model({
18     "optimization_metric": ["f1"],
19     "algorithm": ["dnn"],
20     "name": "anomaly_detection",
21     "data_loader": wrapper_func })
22
23 # Load platform
24 platform = Platforms.Taurus()
25 platform.constrain(
26     "performance": {
27         "throughput": 1, # GPkt/s
28         "latency": 500 }, # ns
29     "resources": { "rows": 16, "cols": 16 })
30
31 # Schedule model and generate code
32 platform.schedule(model_spec)
33 homunculus.generate(platform)
```

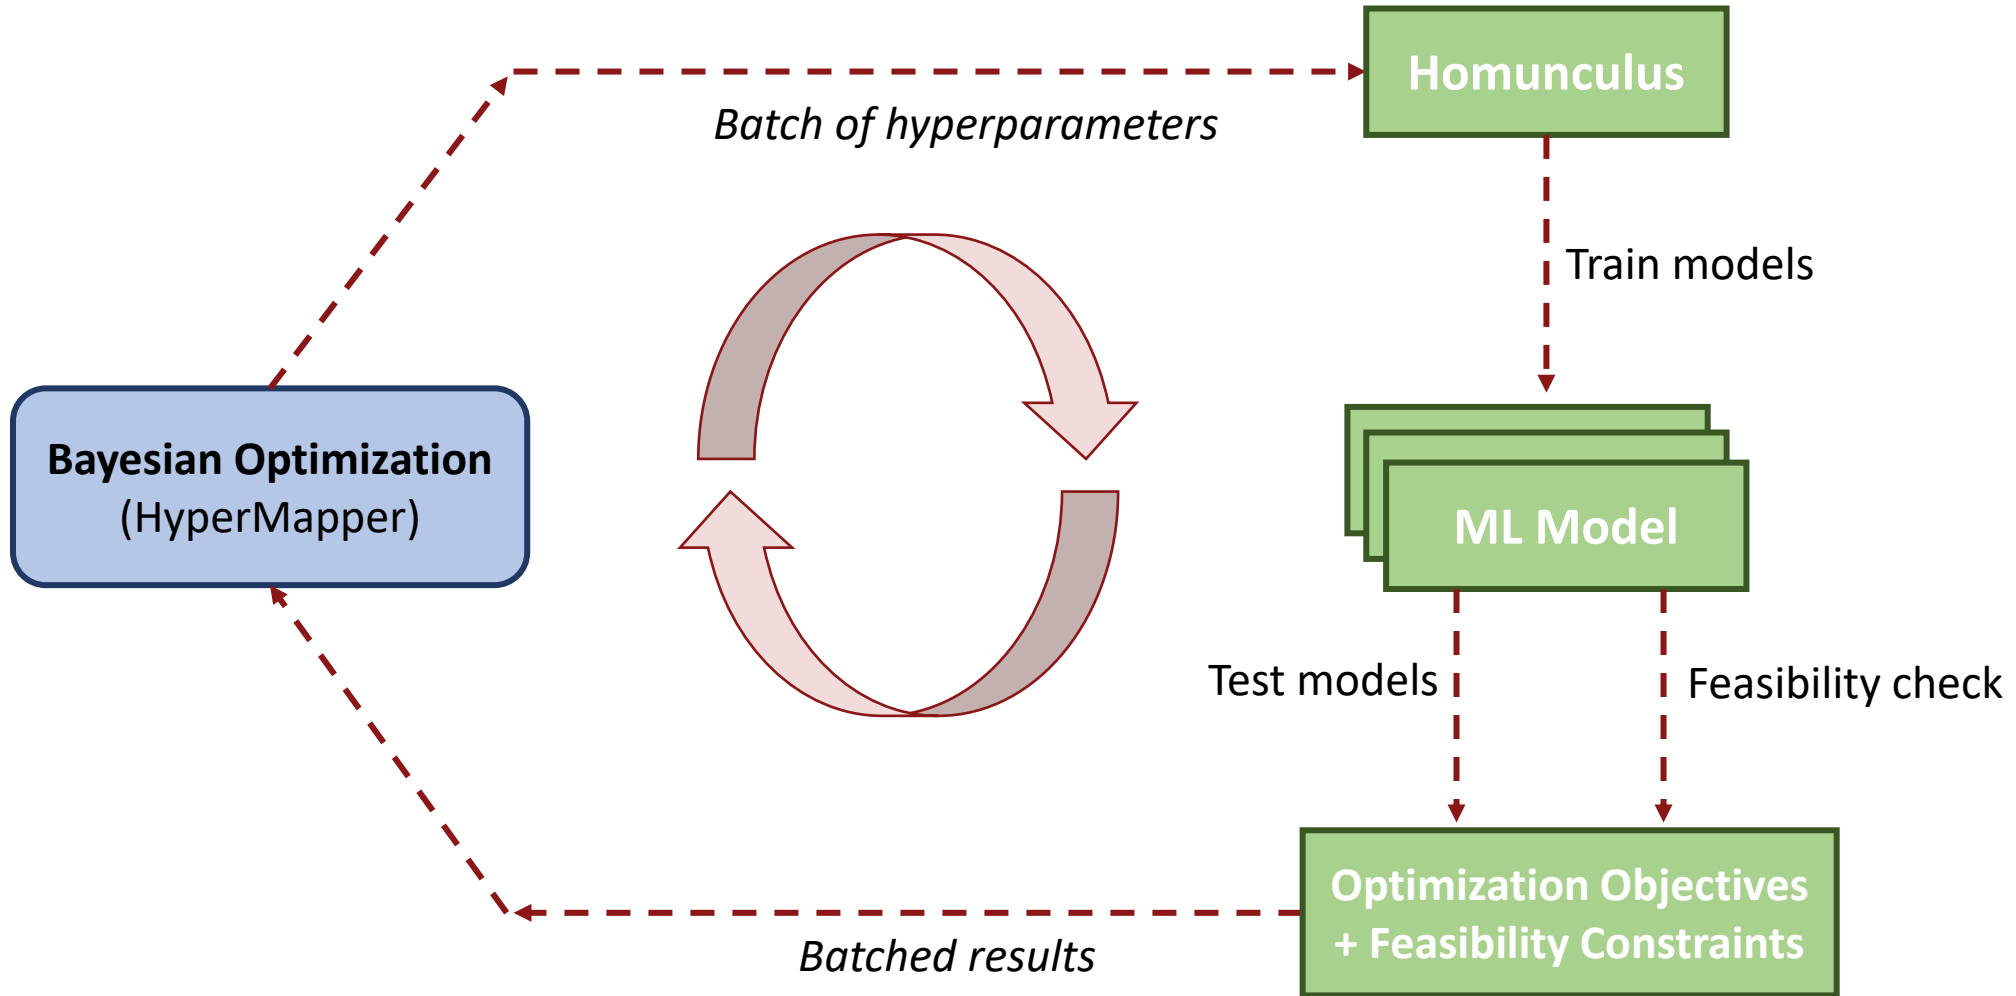
← *Generate binary*



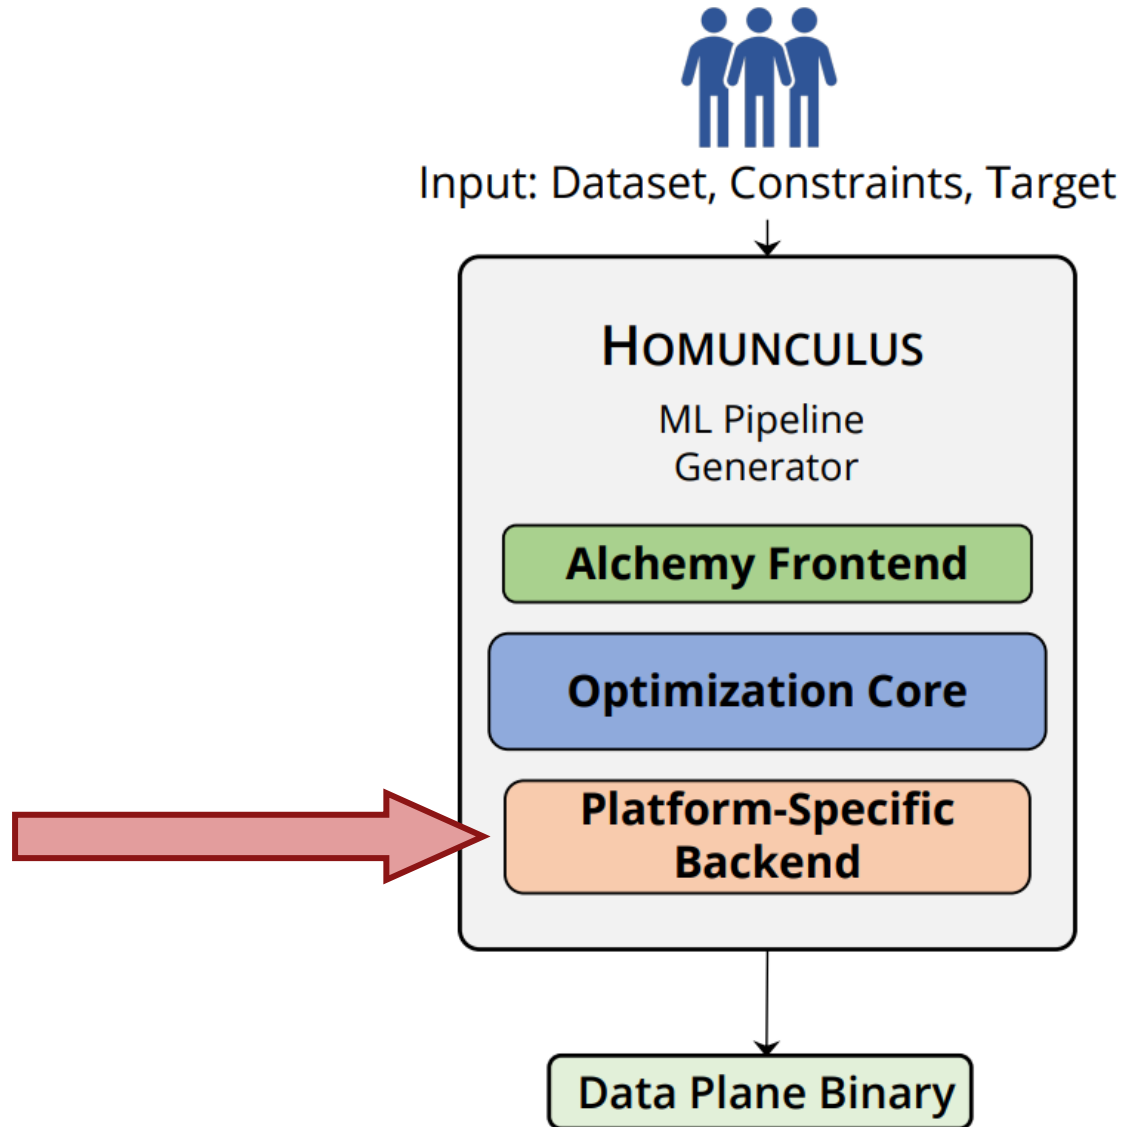
# Homunculus Optimization Core



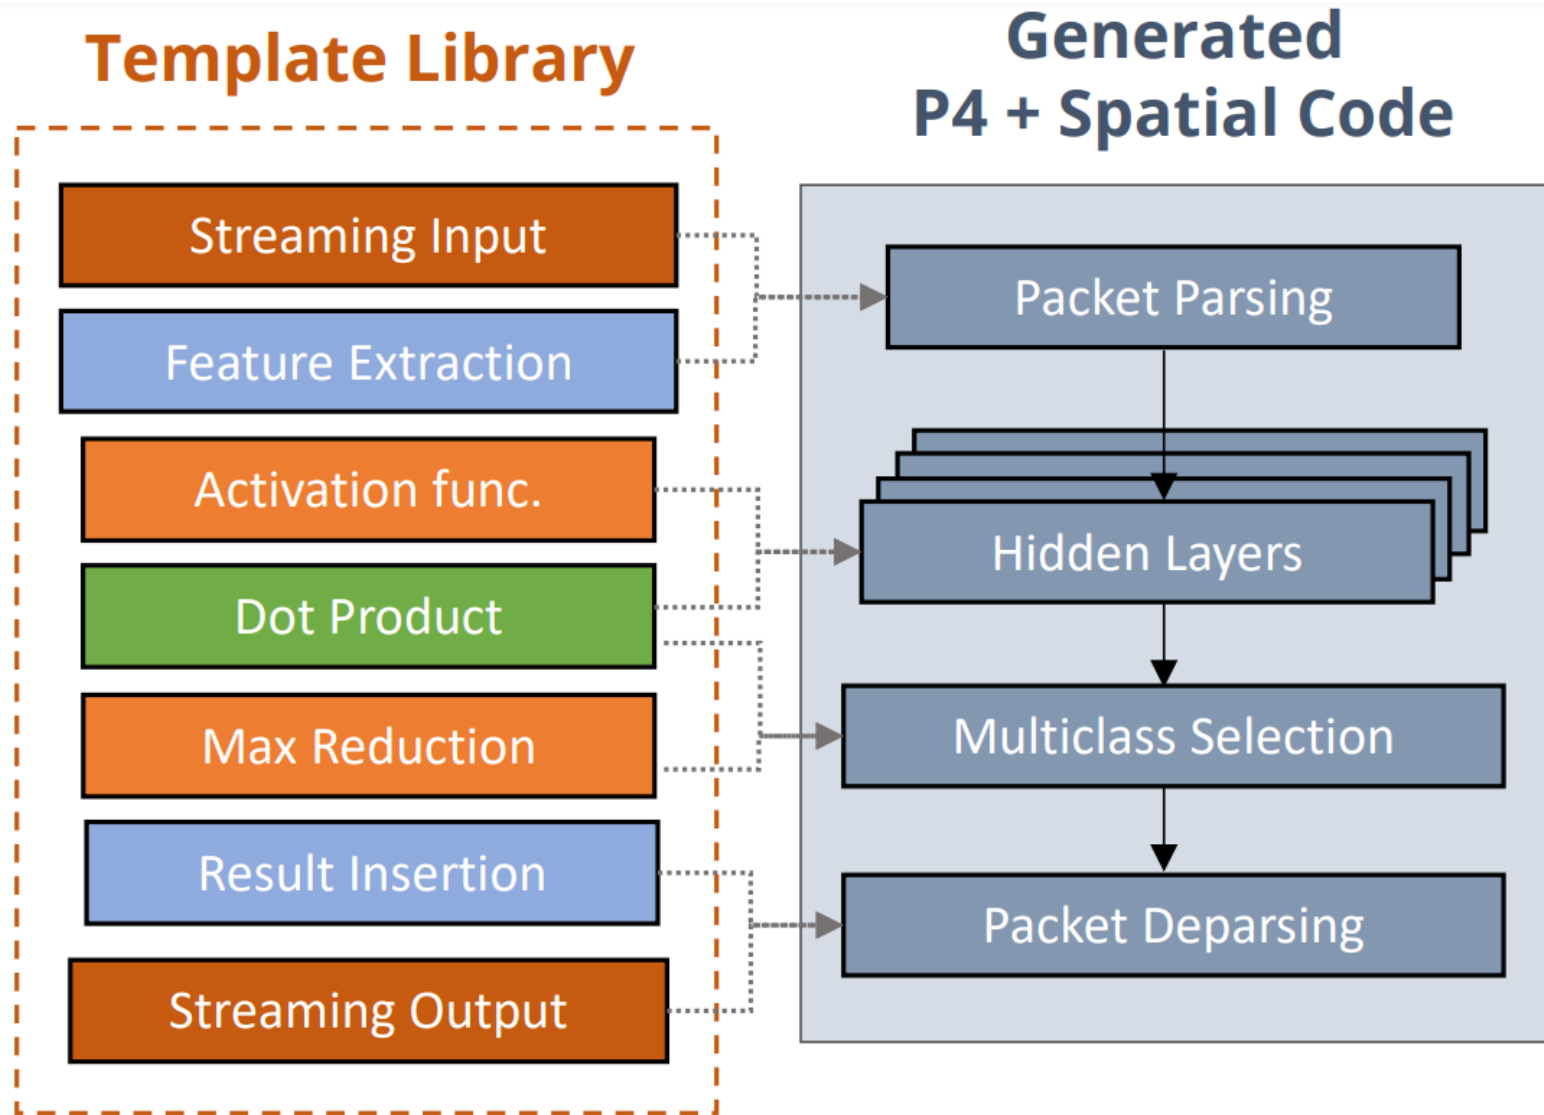
# Homunculus Optimization Core



# Homunculus Backend



# Homunculus Backend Template Library



# Homunculus results

- F1 Score improvements in anomaly detection (**AD**), traffic classification (**TC**), botnet detection (**BD**)

App	Features	# NN Param	F1 Score
Base-AD	7	203	71.10
Hom-AD	7	254	<b>83.10</b>
Base-TC	7	275	61.04
Hom-TC	7	370	<b>68.75</b>
Base-BD	30	662	77.0
Hom-BD	30	501	<b>79.8</b>

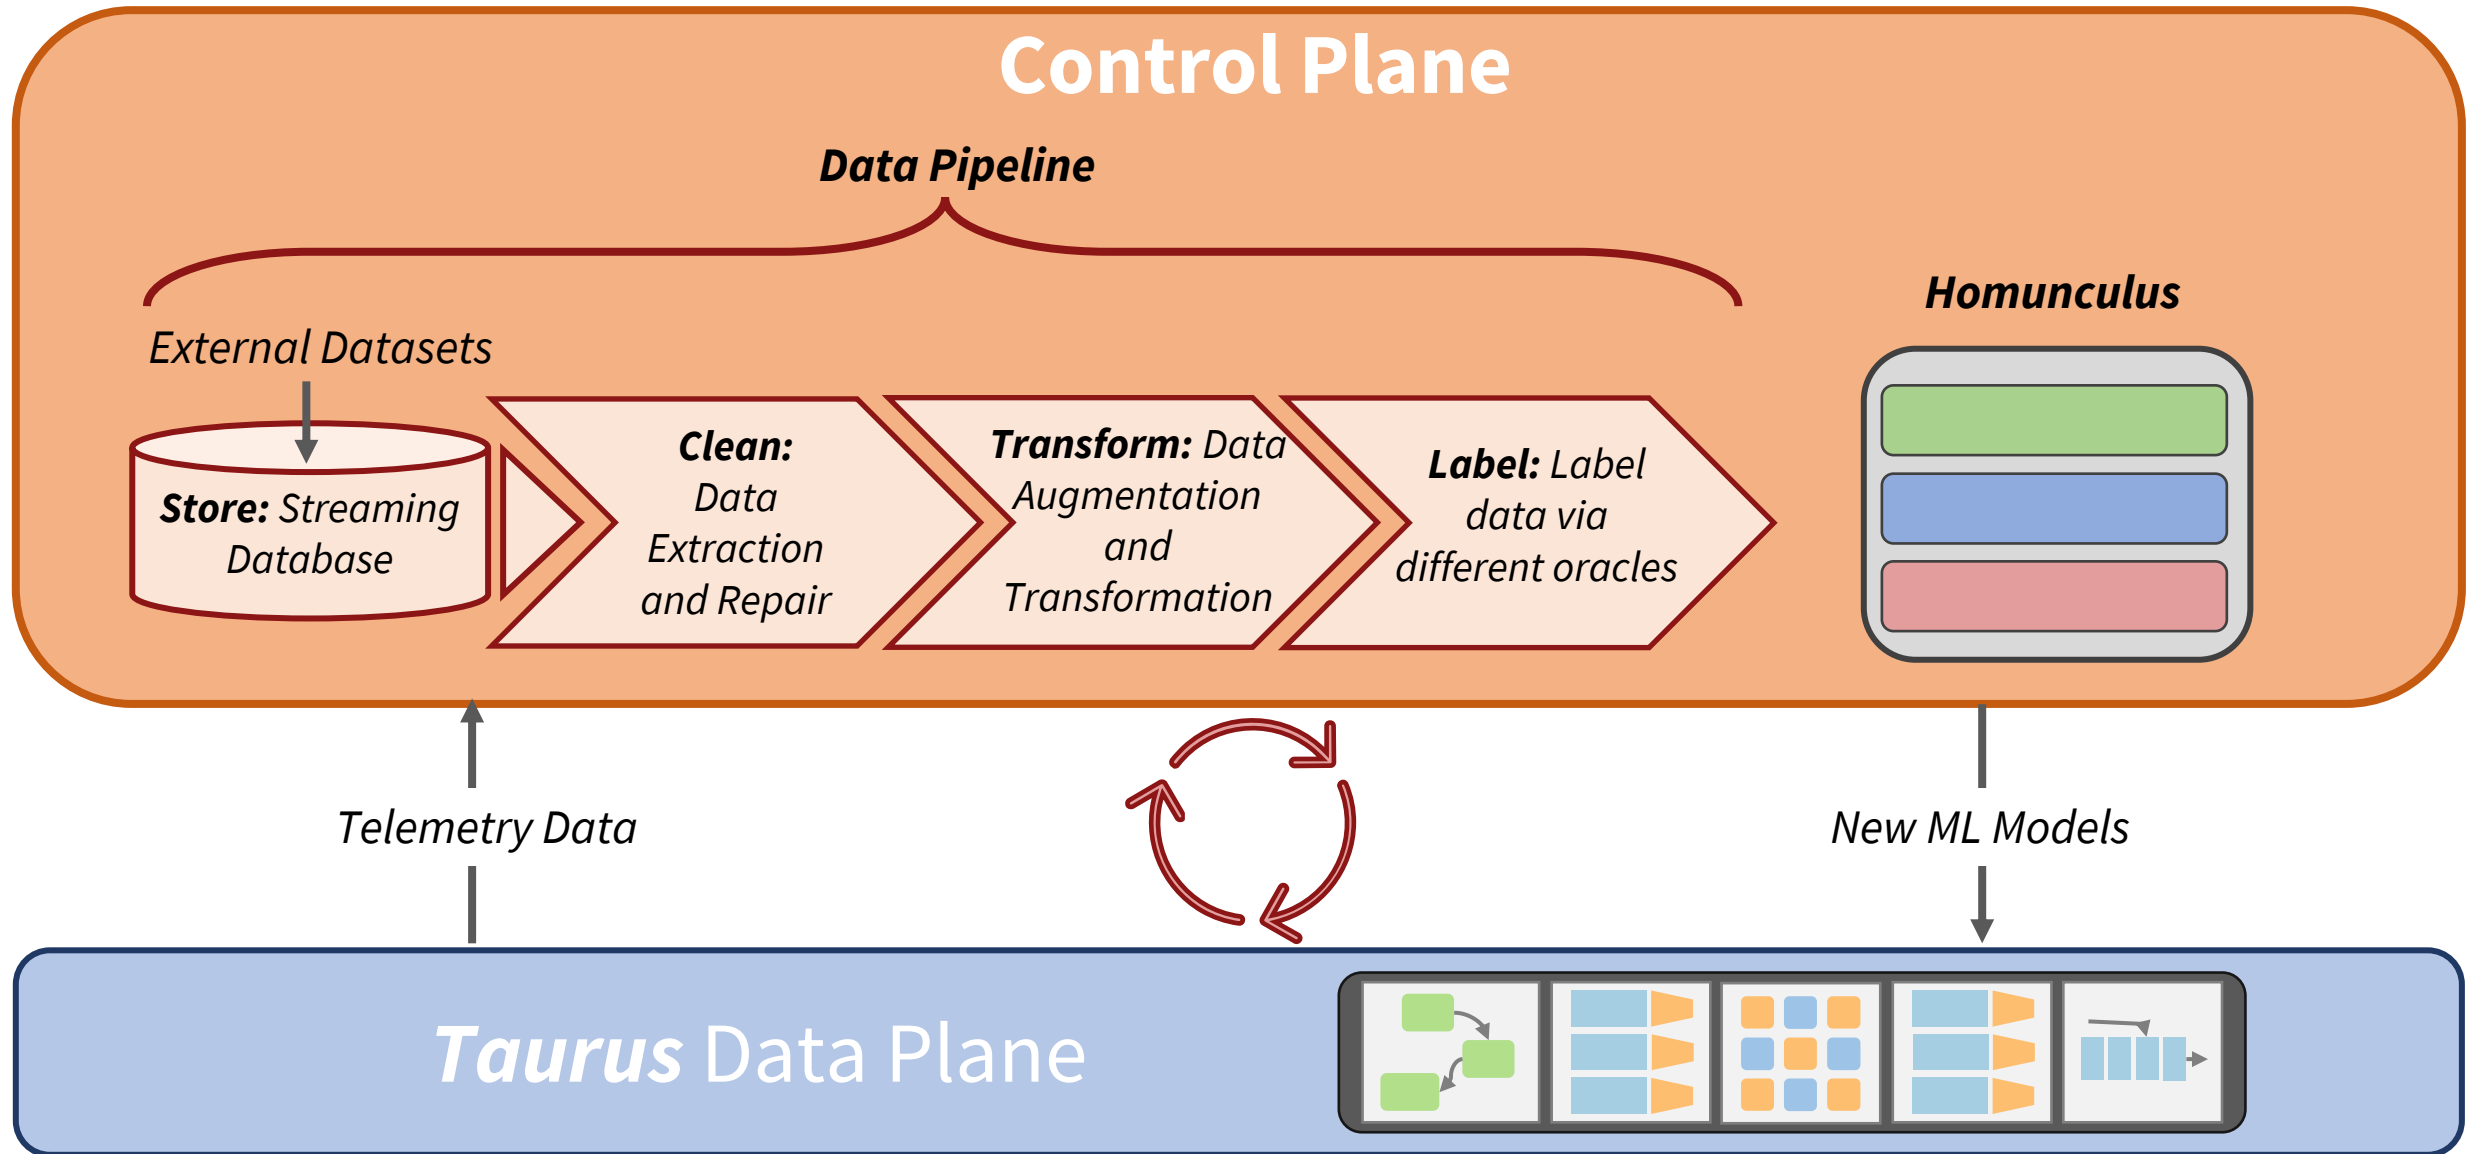
# Homunculus makes it easy to generate high-quality models

- Provides a **high-level interface** for users to express intent
- Uses **network and resource constraints** to traverse AutoML search space
- **Generate binaries** for different dataplane devices with optimized ML models

***Homunculus: Auto-Generating Efficient Data-Plane ML Pipelines for Datacenter Networks*** <https://arxiv.org/pdf/2206.05592>

***How do we supply data to Homunculus?***

# Data Pipeline supplies high quality data to Homunculus





# Questions?

Tushar Swamy

[tswamy@stanford.edu](mailto:tswamy@stanford.edu)

Taurus: <https://dl.acm.org/doi/10.1145/3503222.3507726>




*Enable ML inference  
in the dataplane*



*Run line-rate, per-packet  
operation*

Homunculus: <https://arxiv.org/pdf/2206.05592>



*Allows automated  
construction of ML models*



*More accurate than  
hand-tuned models*

Try it out! <https://gitlab.com/dataplane-ai/taurus>