# Measuring and Forecasting Casualties of Political Violence

Benjamin J. Radford

November 10, 2024

University of North Carolina at Charlotte

# Benjamin Radford

- Political methodologist
    - Machine learning
    - Bayesian statistics
    - Forecasting
    - Measurement
    - Political conflict & violence
- Assistant Professor at UNC Charlotte
    - Political Science & Public Administration
    - Public Policy PhD Program
    - School of Data Science
- Data Scientist  various companies
    - Enhanced Attribution (Defense Advanced Research Projects Agency)
    - Network Defense (DARPA)
    - INCAS (DARPA)

1. Forecasting Battle Deaths
2. Forecasting Protests in Hong Kong
3. Measuring Battle Deaths
4. Conclusion

Companion Webpage (benradford.com)

# Forecasting Battle Deaths

# High Resolution Conflict Forecasting

INTERNATIONAL INTERACTIONS
2022, VOL. 48, NO. 4, 739–758
https://doi.org/10.1080/03050629.2022.2031182

Routledge
Taylor & Francis Group



## High resolution conflict forecasting with spatial convolutions and long short-term memory

Benjamin J. Radford

University of North Carolina at Charlotte

**ABSTRACT**

The 2020 Violence Early Warning System (ViEWS) Prediction Competition challenged participants to produce predictive models of violent political conflict at high spatial and temporal resolutions. This paper presents a convolutional long short-term memory (ConvLSTM) recurrent neural network capable of forecasting the log change in battle-related deaths resulting from state-based armed conflict at the PRIO-GRID cell-month level. The ConvLSTM outperforms the benchmark model provided by the ViEWS team and performs comparably to the best models submitted to the competition. In addition to providing a technical description of the ConvLSTM, I evalu-

**KEYWORDS**
Forecasting; machine learning; neural networks; political conflict

# The Question

Can we predict changes in state-based violence?

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
    - War is in the Error Term [Gartzke, 1999]
    - Forecasting to avoid overfitting observational models [Ward et al., 2013]
    - DARPA's Integrated Crisis Early Warning System
    - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
    - PRIO's Violence Early Warning System (ViEWS)

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
    - War is in the Error Term [Gartzke, 1999]
    - Forecasting to avoid overfitting observational models [Ward et al., 2013]
    - DARPA's Integrated Crisis Early Warning System
    - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
    - PRIO's Violence Early Warning System (ViEWS)

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
  - War is in the Error Term [Gartzke, 1999]
  - Forecasting to avoid overfitting observational models [Ward et al., 2013]
  - DARPA's Integrated Crisis Early Warning System
  - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
  - PRIO's Violence Early Warning System (ViEWS)

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
  - War is in the Error Term [Gartzke, 1999]
  - Forecasting to avoid overfitting observational models [Ward et al., 2013]
  - DARPA's Integrated Crisis Early Warning System
  - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
  - PRIO's Violence Early Warning System (ViEWS)

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
    - War is in the Error Term [Gartzke, 1999]
    - Forecasting to avoid overfitting observational models [Ward et al., 2013]
    - DARPA's Integrated Crisis Early Warning System
    - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
    - PRIO's Violence Early Warning System (ViEWS)

# Forecasting Conflict is Science

- Forecasting is about making *falsifiable* predictions.
- If we can forecast conflict, we can (perhaps) mitigate violence.
- Forecasting conflict has a long history in the social sciences:
    - War is in the Error Term [Gartzke, 1999]
    - Forecasting to avoid overfitting observational models [Ward et al., 2013]
    - DARPA's Integrated Crisis Early Warning System
    - ...Lots of great research [Muchlinski et al., 2016, Colaresi and Mahmood, 2017]
    - PRIO's Violence Early Warning System (ViEWS)

Uppsala University and the Peace Research Institute of Oslo (PRIO)

The Violence & Impacts Early-Warning System (VIEWS)

ViEWS Prediction Competition (2020)

# What is a Shared Task?

## Shared Task

A collaborative research effort in which teams compete to accomplish a task given shared data and evaluation metrics.

## Components of Shared Tasks

- Objective or "task"
- Data
  - Training data ← fit model to this
  - Validation data ← use to perform model search or tuning
  - Evaluation/Test data ← unobserved until eval time
- Target metrics (e.g., MSE, Accuracy, ...)

# Shared Tasks

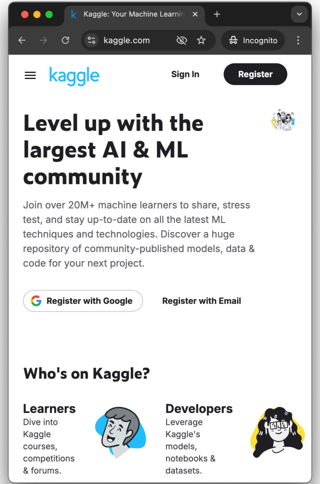


- Netflix Prize (2009)
  - Build a better film recommendation engine
- Dozens of NLP tasks for LLMs
  - **The Abstract Reasoning Challenge** [Chollet, 2019]
  - **General Language Understanding Evaluation** [Wang et al., 2018]
- Basically every other possible task:
  - **Passenger Screening Algorithm Challenge** $1,500,000. Dept. of Homeland Security
  - **Home Value Prediction** $1,200,000. Zillow
  - **Deepfake Detection Challenge** $1,000,000. Amazon, Meta, Microsoft

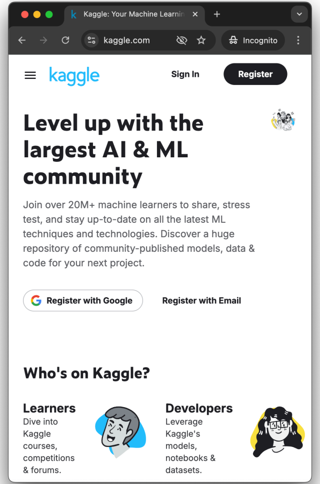


- Netflix Prize (2009)
  - Build a better film recommendation engine
- Dozens of NLP tasks for LLMs
  - **The Abstract Reasoning Challenge** [Chollet, 2019]
  - **General Language Understanding Evaluation** [Wang et al., 2018]
- Basically every other possible task:
  - **Passenger Screening Algorithm Challenge** $1,500,000. Dept. of Homeland Security
  - **Home Value Prediction** $1,200,000. Zillow
  - **Deepfake Detection Challenge** $1,000,000. Amazon, Meta, Microsoft

# Shared Tasks

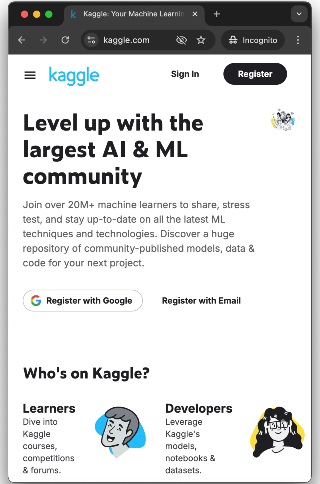


- Netflix Prize (2009)
  - Build a better film recommendation engine
- Dozens of NLP tasks for LLMs
  - **The Abstract Reasoning Challenge** [Chollet, 2019]
  - **General Language Understanding Evaluation** [Wang et al., 2018]
- Basically every other possible task:
  - **Passenger Screening Algorithm Challenge** $1,500,000. Dept. of Homeland Security
  - **Home Value Prediction** $1,200,000. Zillow
  - **Deepfake Detection Challenge** $1,000,000. Amazon, Meta, Microsoft

# ViEWS Goal (Country-Month)

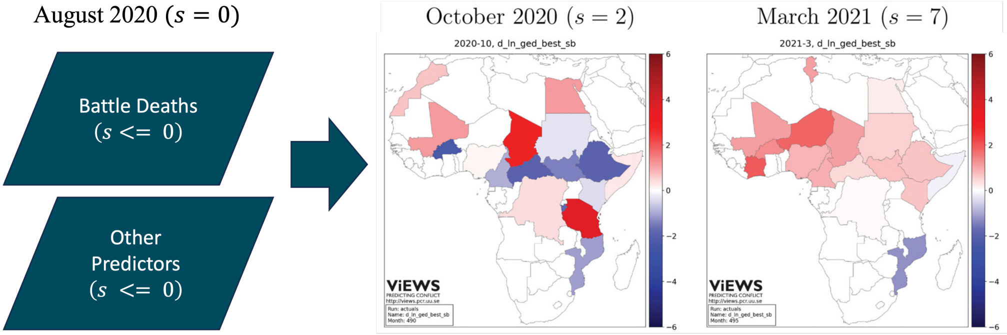


**Figure 1:** ViEWS Competition illustration (*cm*) [Hegre et al., 2022]

# ViEWS Goal (PRIO Grid-Month)

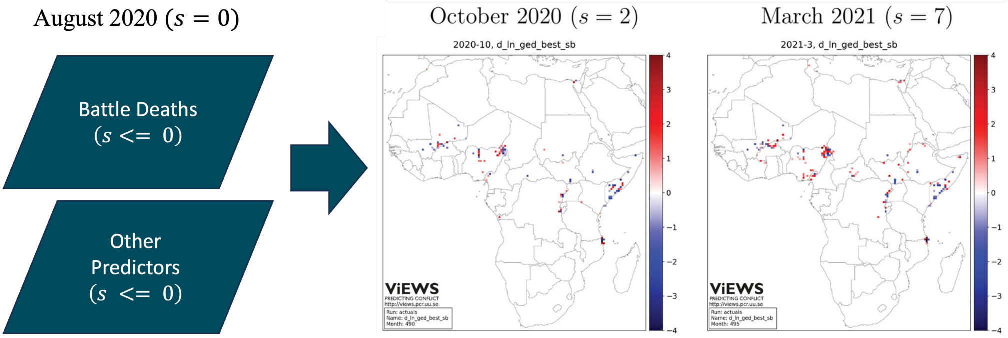


**Figure 2:** ViEWS Competition illustration (*pgm*) [Hegre et al., 2022]

## PRIO Grid-Month

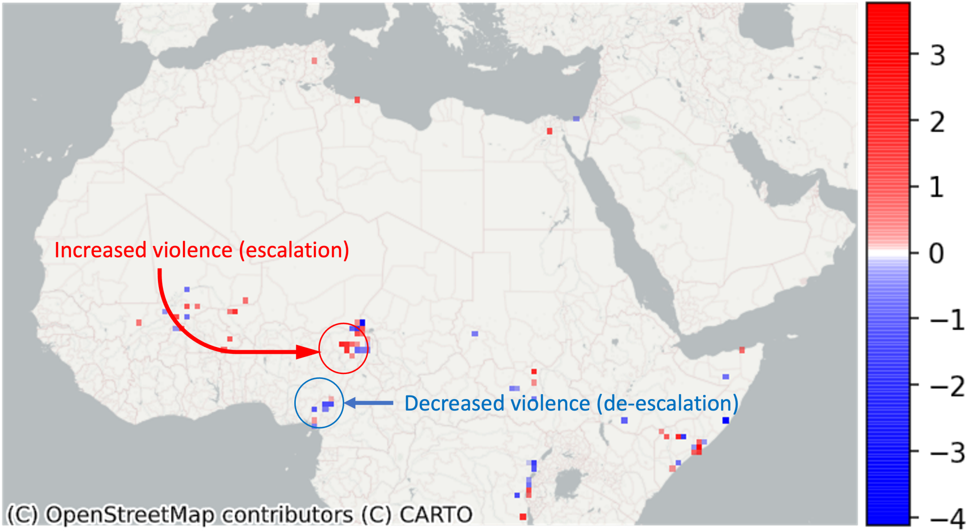

### The Target (DV)

Change in *Battle Related Deaths* / grid-cell month.
UCDP-GED [Sundberg and Melander, 2013]

### The Data

- 1990–2013: Training
- 2014–2016: Validation
- 2017–2019: Testing
- 2020–2021: True out-of-sample

### Evaluation

**Metrics**: MSE, TADDA, MAL, pEMDiv
**Benchmark model**: random forest

## ViEWS Competition

- Make predictions of $\Delta ln(\text{battle deaths} + 1)_{t+s}$
- Resolution: monthly grid cells
- Grid cells: $\sim$2500.0 km$^2$ (one-half degree lat / lon)
- Time frame: 1990–2020

# What Do the Data Look Like?

Let’s consider what our data “look like” to motivate our modeling choices.

# What Does the Target Look Like?

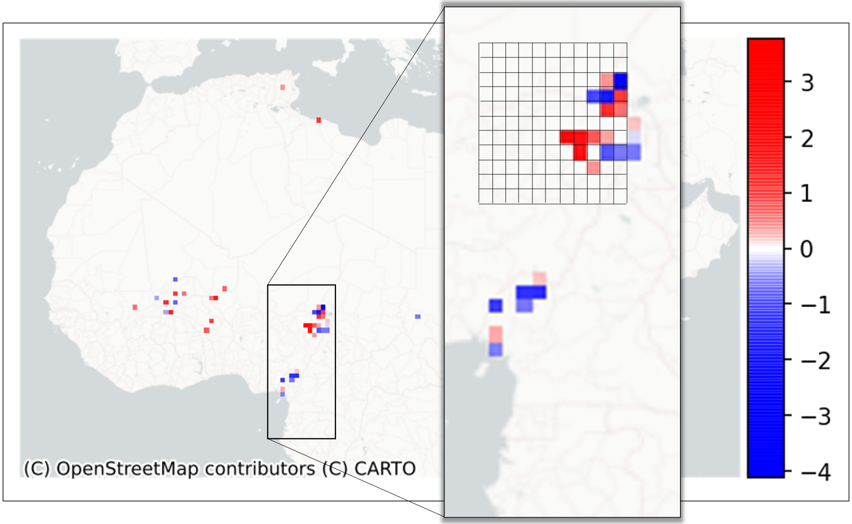


Figure 4: $\Delta ln(\text{battle deaths} + 1)$

| | Variable | Description |
|---|---|---|
| 1 | ln_ged_best_sb | Current $\ln(\text{deaths} + 1)$ |
| 2 | pgd_bdist3 | Border distance (km) |
| 3 | pgd_capdist | Distance to capital (km) |
| 4 | pgd_agri_ih | Agricultural area % |
| 5 | pgd_pop_gpw_sum | Population |
| 6 | pgd_ttime_mean | Travel time to major city |
| 7 | spdist_pgd_diamsec | Diamond resources |
| 8 | pgd_pasture_ih | Pasture area % |
| 9 | pgd_savanna_ih | Savanna area % |
| 10 | pgd_forest_ih | Forest area % |
| 11 | pgd_urban_ih | Urban area % |
| 12 | pgd_barren_ih | Barren area % |
| 13 | pgd_gcp_mer | Gross cell product (USD) |

# Features

| | Variable | Description |
|---|---|---|
| 1 | ln_ged_best_sb | Current $\ln(\text{deaths} + 1)$ |
| 2 | pgd_bdist3 | Border distance (km) |
| 3 | pgd_capdist | Distance to capital (km) |
| 4 | pgd_agri_ih | Agricultural area % |
| 5 | pgd_pop_gpw_sum | Population |
| 6 | pgd_ttime_mean | Travel time to major city |
| 7 | spdist_pgd_diamsec | Diamond resources |
| 8 | pgd_pasture_ih | Pasture area % |
| 9 | pgd_savanna_ih | Savanna area % |
| 10 | pgd_forest_ih | Forest area % |
| 11 | pgd_urban_ih | Urban area % |
| 12 | pgd_barren_ih | Barren area % |
| 13 | pgd_gcp_mer | Gross cell product (USD) |

# Feature Maps

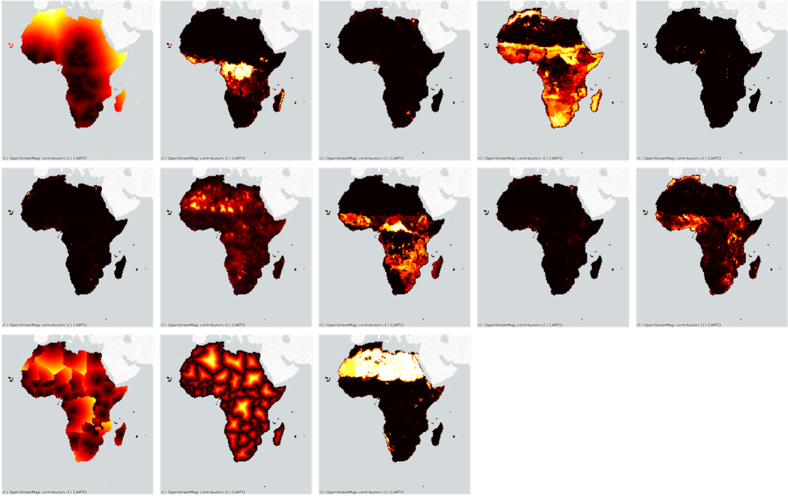

# Feature per Month

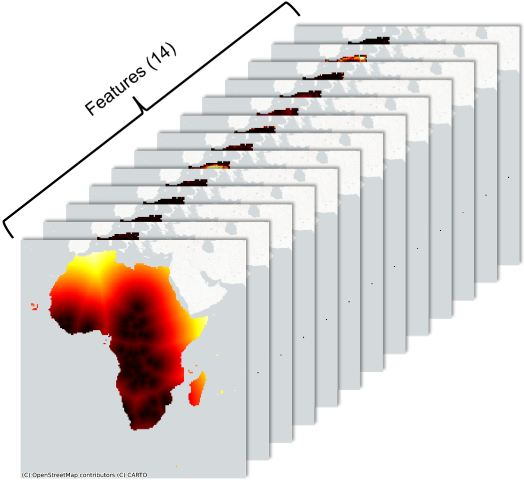

# Features Over Time

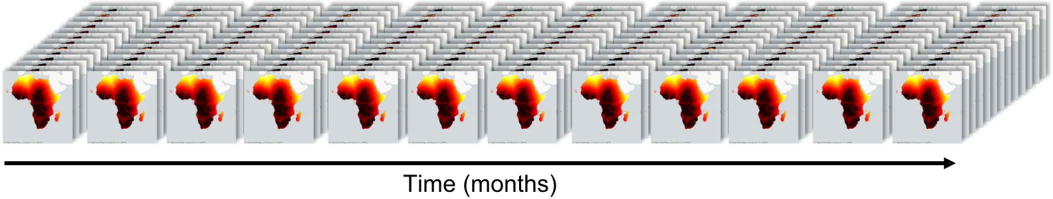

# Reshape

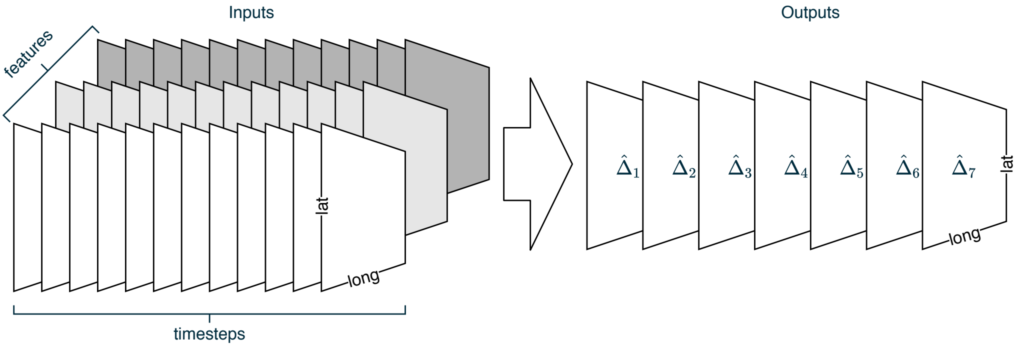

Inputs
features
lat
long
timesteps
Outputs
$\hat{\Delta}_1$
$\hat{\Delta}_2$
$\hat{\Delta}_3$
$\hat{\Delta}_4$
$\hat{\Delta}_5$
$\hat{\Delta}_6$
$\hat{\Delta}_7$
lat
long

# Input and Output Tensors

**Input Size (single sample)**

$$(12 \times 178 \times 169 \times 14) = \mathbf{5,053,776}$$

**Output Size (single sample)**

$$(7 \times 178 \times 169 \times 1) = \mathbf{210,574}$$

**Training Set Size**

$$5,053,776/12 \times 270 = \mathbf{113,709,960}$$

# Model Requirements

## Attributes We Want

- Spatial effects
- Temporal effects
- Additional features / covariates

## This Looks Like...

Next frame prediction for generating videos [Lotter et al., 2017]

## The Model

Convolutional Long Short-Term Memory Neural Network (ConvLSTM)

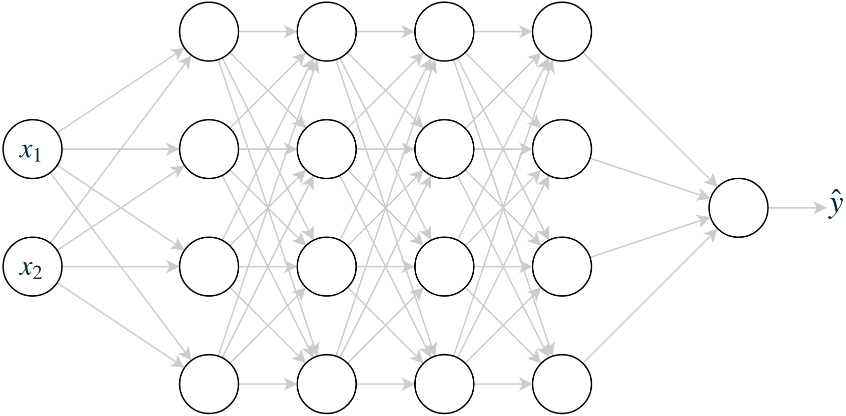

$x_1$
$x_2$
$\hat{y}$

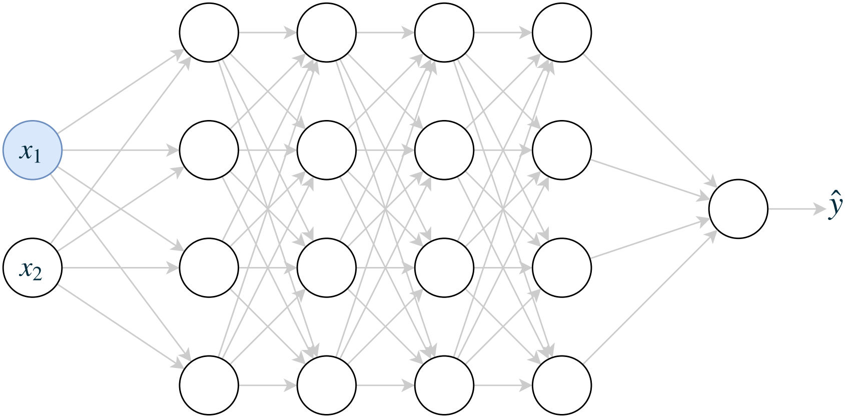
$x_1$
$x_2$
$\hat{y}$

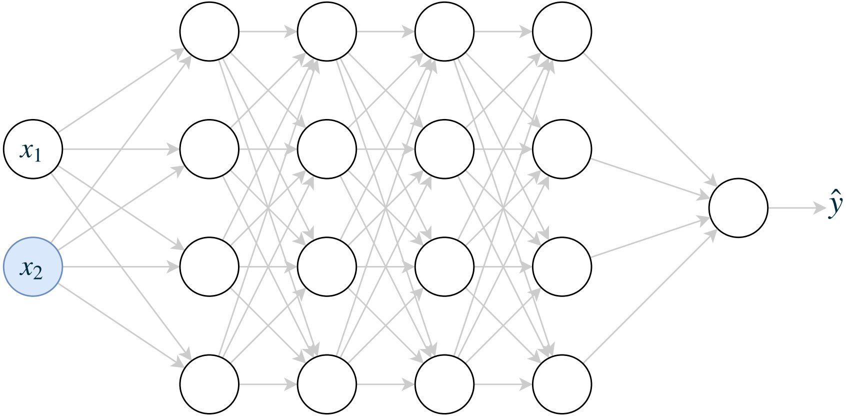
$x_1$
$x_2$
$\hat{y}$

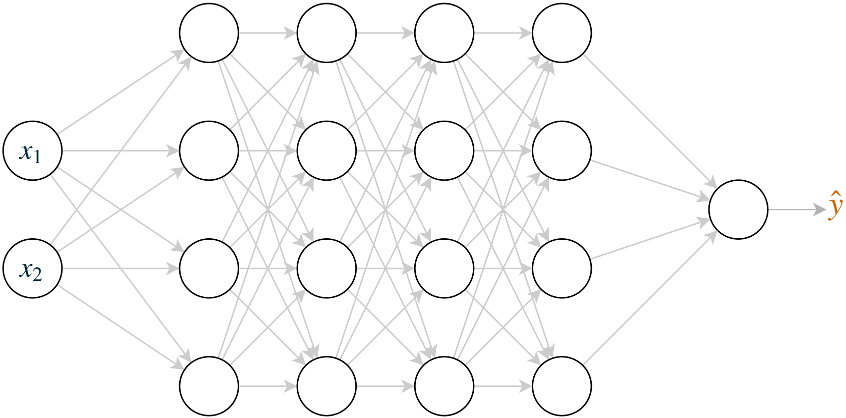


This is $\hat{y}$, our predicted value for $y$.

It is the output of our neural network.

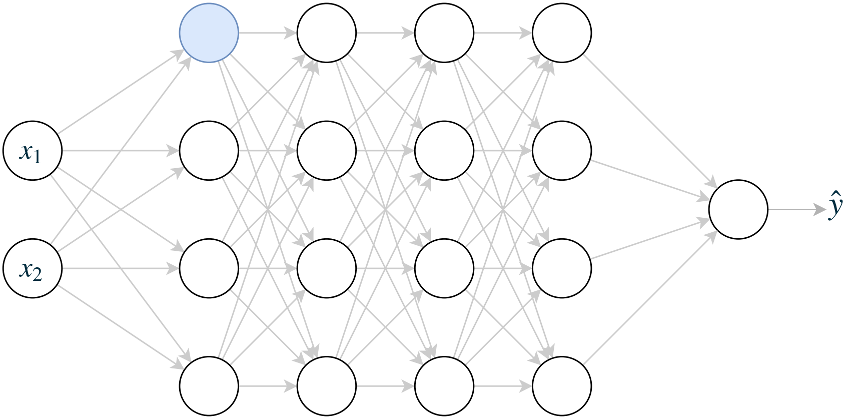


This is a neuron.

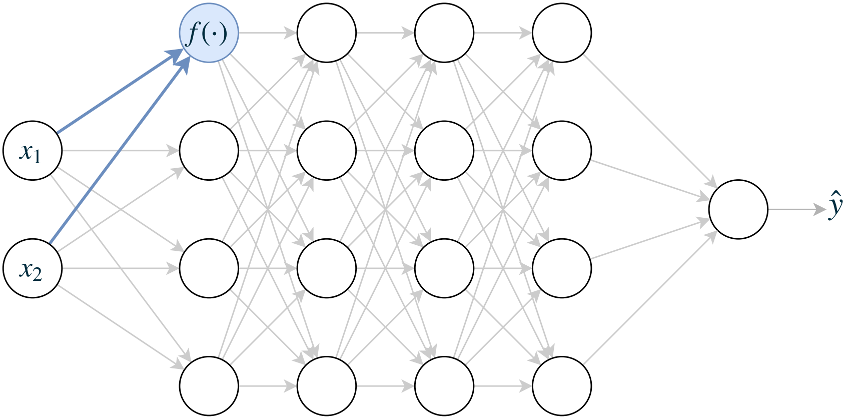


A neuron is a function.

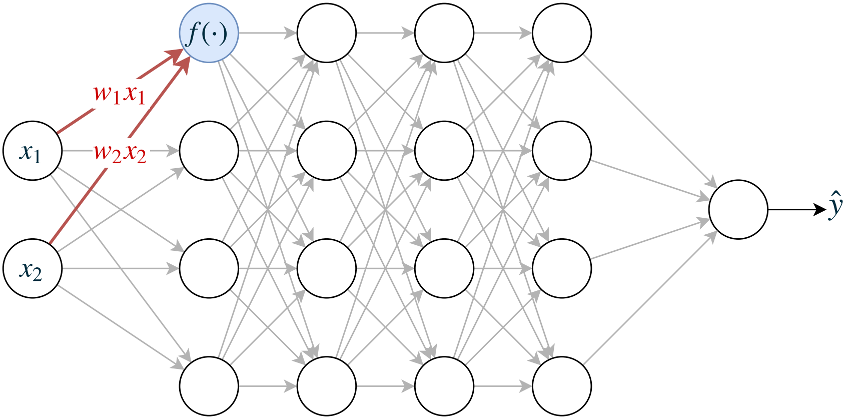


Our neuron's input is $(w_1 \times x_1) + (w_2 \times x_2)$.

$w_1$ and $w_2$ are weights, or coefficients.

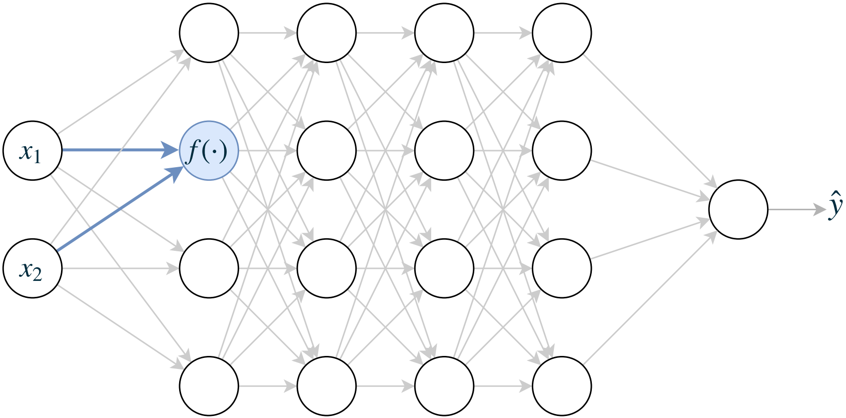


This is another neuron.

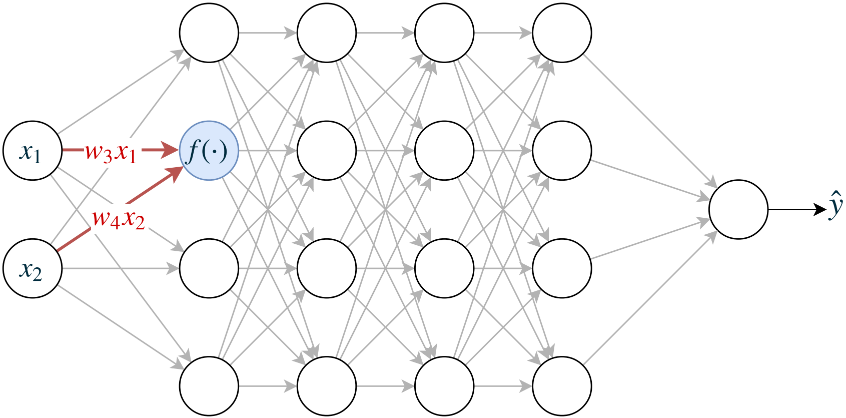


It's value is a function of $(w_3 \times x_1) + (w_4 \times x_2)$.

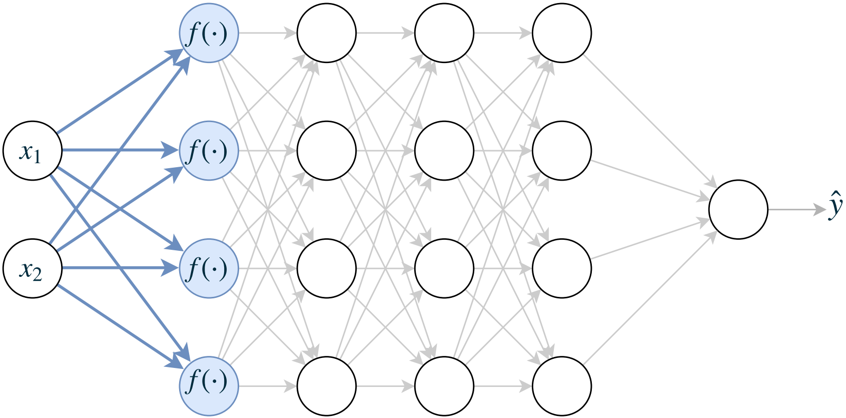


A **layer** is a bunch of neurons shared inputs and different weights.
Every arrow above is a unique weight (or coefficient).

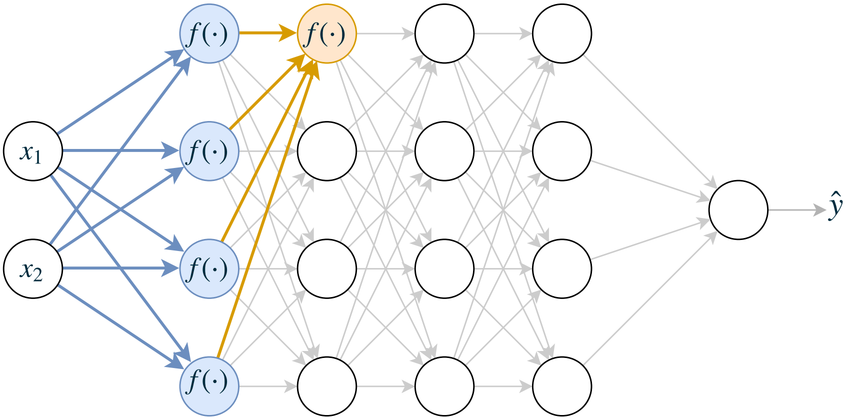


The inputs to **this neuron** are the outputs of the **previous layer**.

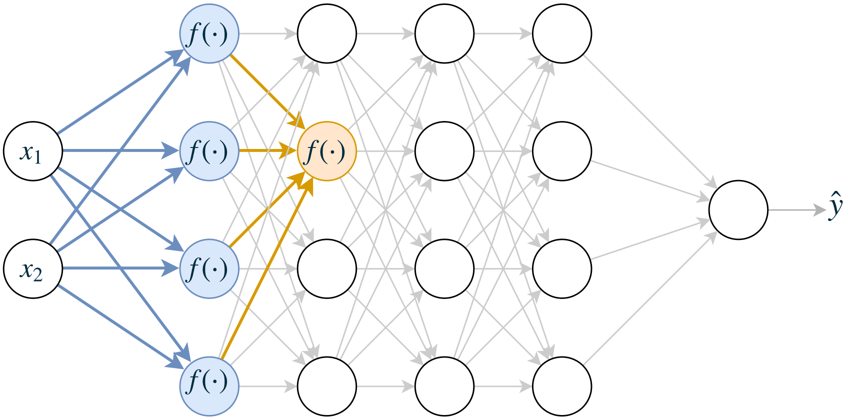


The inputs to **this neuron** are the outputs of the **previous layer**.

# Neural Networks

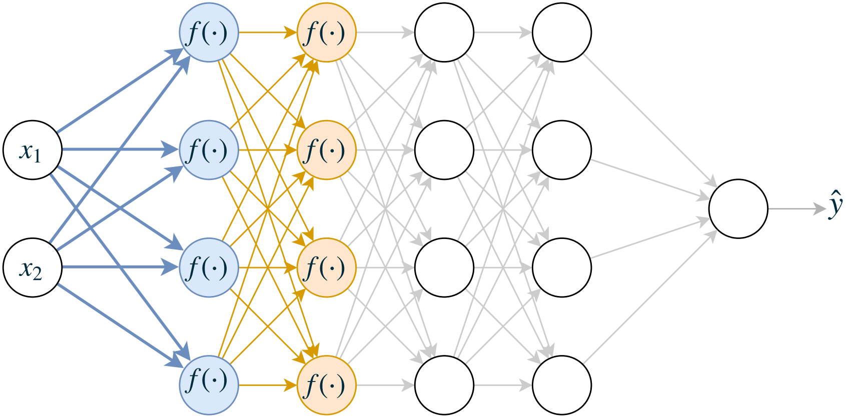

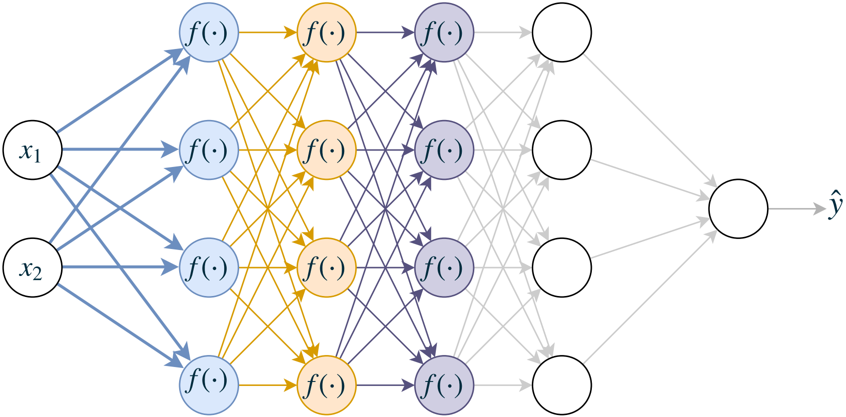
$x_1$
$x_2$
$f(\cdot)$
$\hat{y}$

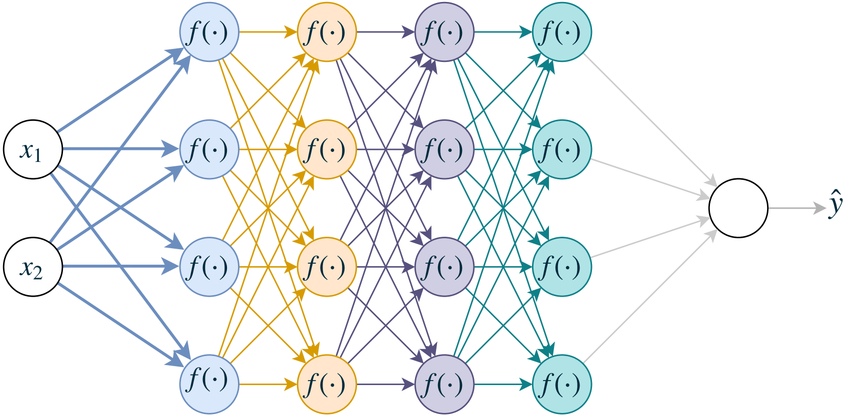


These are hidden layers.

# Neural Networks

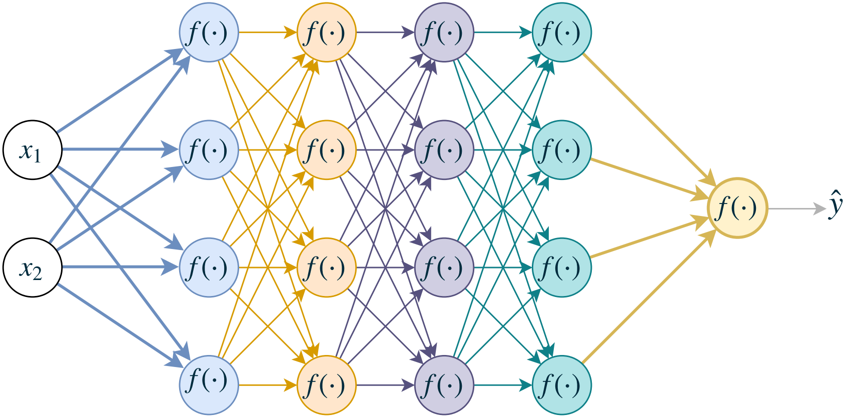

# Convolutional

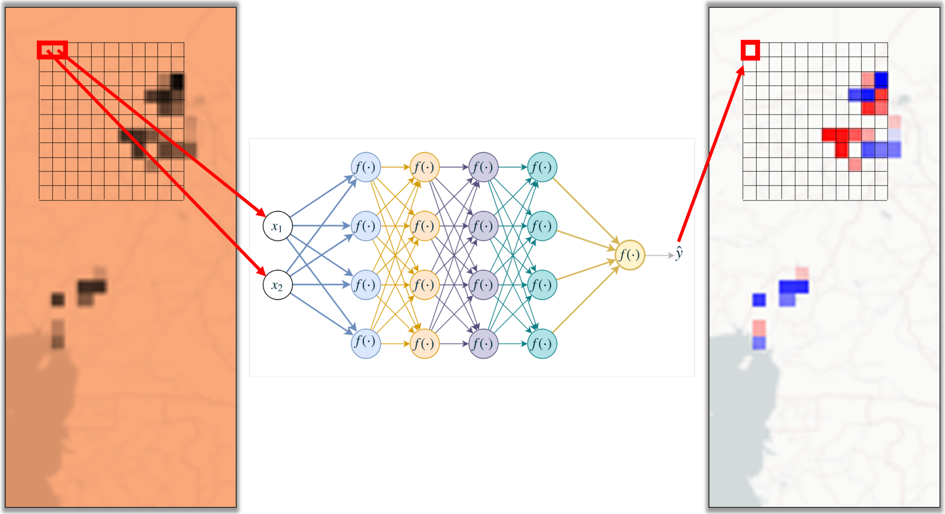

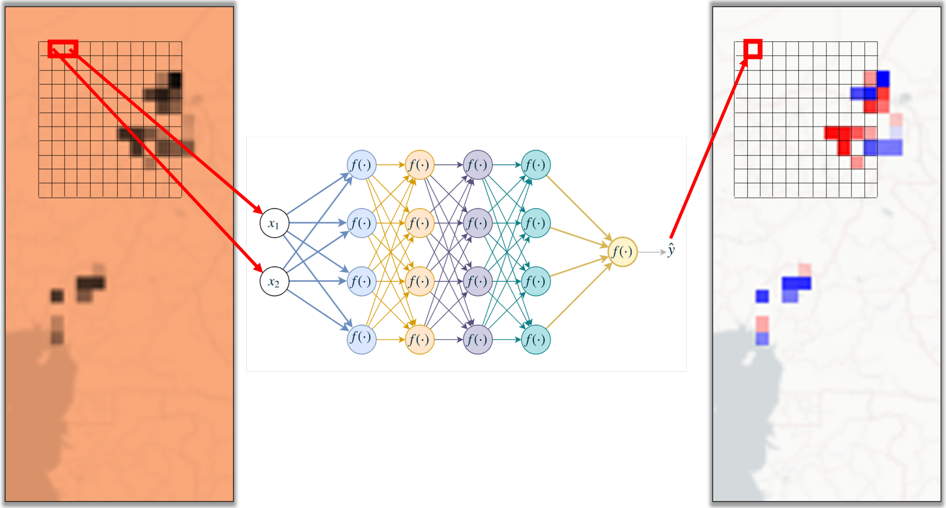
$x_1$
$x_2$
$f(\cdot)$
$\hat{y}$

# Convolutional

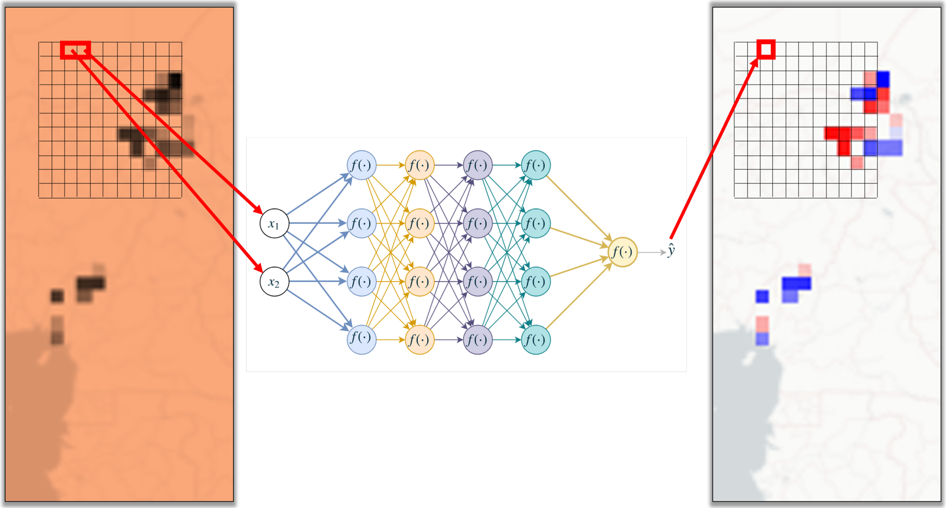

# Convolutional

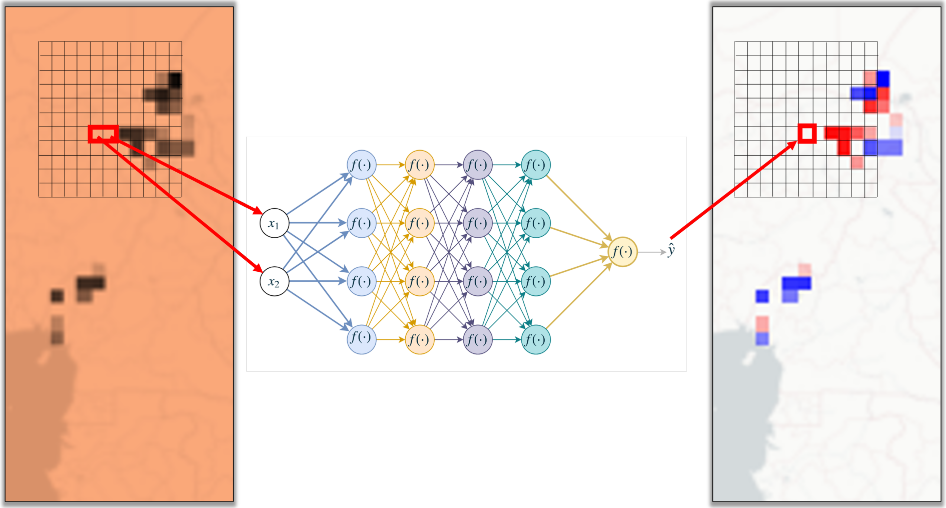

# Convolutional

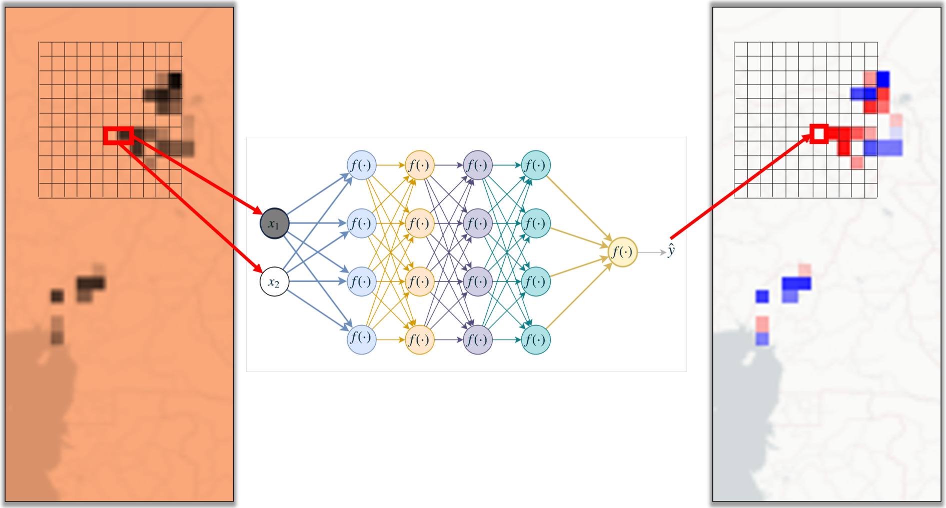

# Convolutional

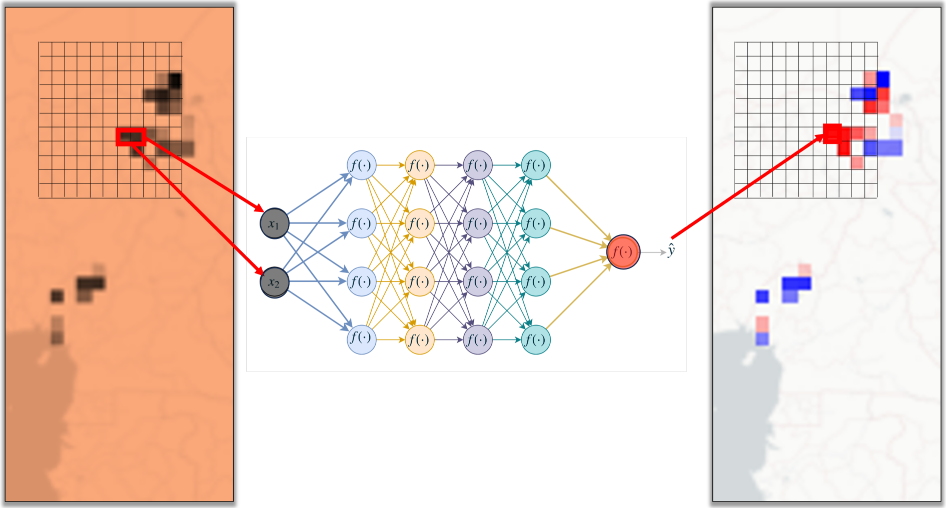

# LSTM / Recurrent Neural Network

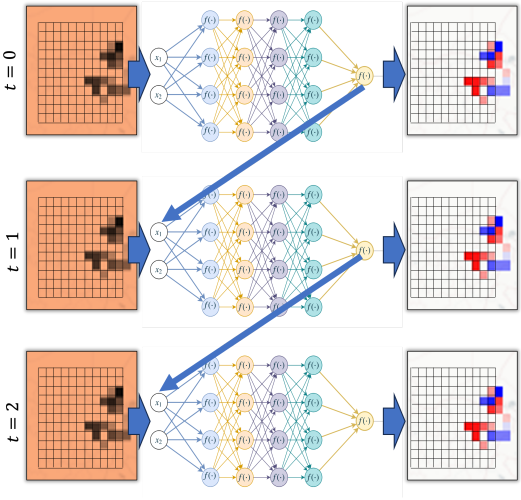

# RESHAPE

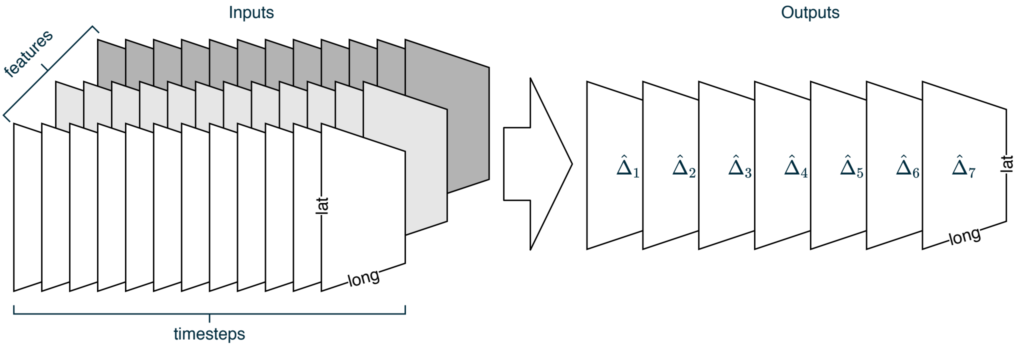

Inputs
Outputs
features
timesteps
lat
long
$\hat{\Delta}_1$
$\hat{\Delta}_2$
$\hat{\Delta}_3$
$\hat{\Delta}_4$
$\hat{\Delta}_5$
$\hat{\Delta}_6$
$\hat{\Delta}_7$
lat
long

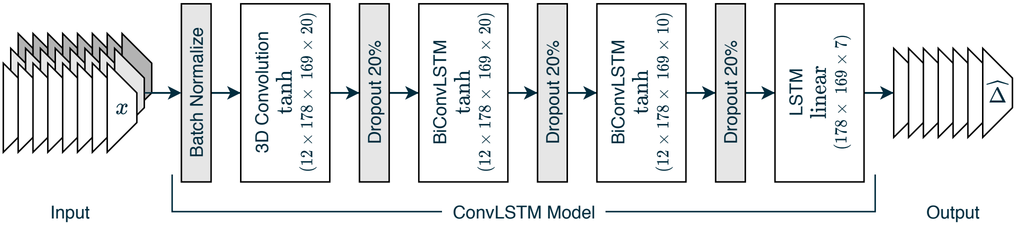

$x$
Batch Normalize
3D Convolution
tanh
$(12 \times 178 \times 169 \times 20)$
Dropout 20%
BiConvLSTM
tanh
$(12 \times 178 \times 169 \times 20)$
Dropout 20%
BiConvLSTM
tanh
$(12 \times 178 \times 169 \times 10)$
Dropout 20%
LSTM
linear
$(178 \times 169 \times 7)$
$\widehat{\Delta}$
Input
ConvLSTM Model
Output

# Model Details

- **Parameters**: 281,016
- **Loss**: MSE
- **Optimizer**: RMSprop
- **Batch Size**: 8
- **Epochs**: 75


**Training time**: about 1.5 hours

# Max Predictions in Test Set (+2 Months)

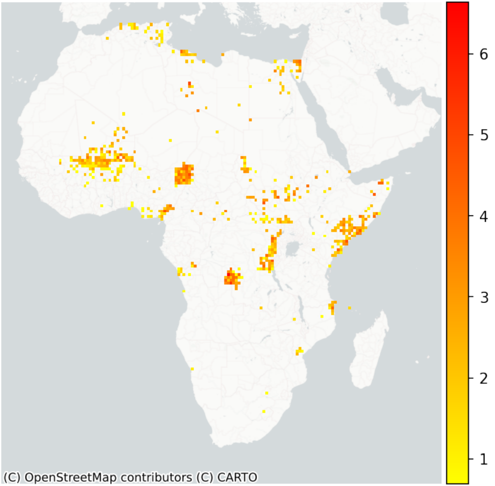


**Figure 5:** Observed Max

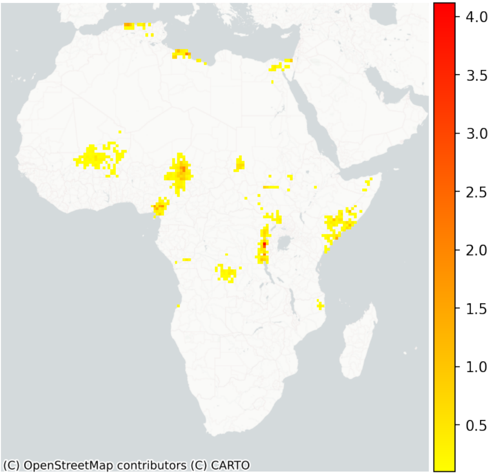


**Figure 6:** Predicted Max

# Min Predictions in Test Set (+2 Months)

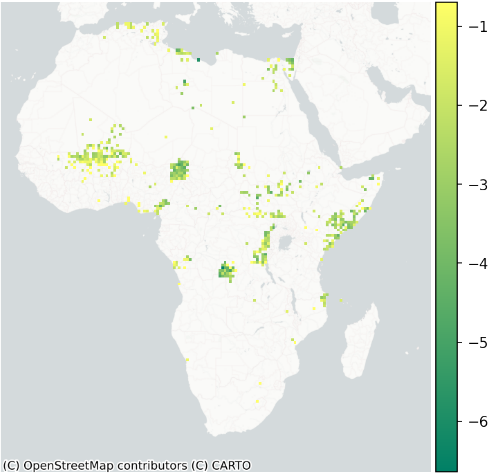


Figure 7: Observed Min

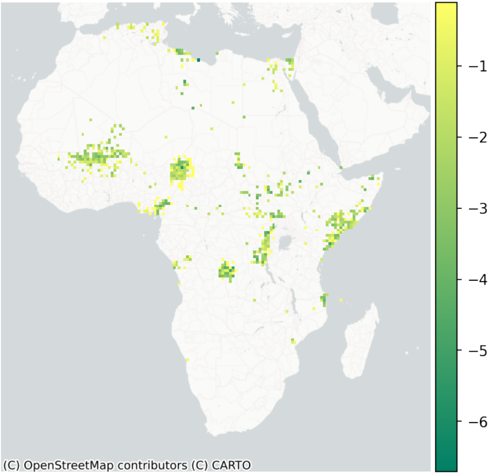


Figure 8: Predicted Min

# Actual: December 2018 (+2 Months)

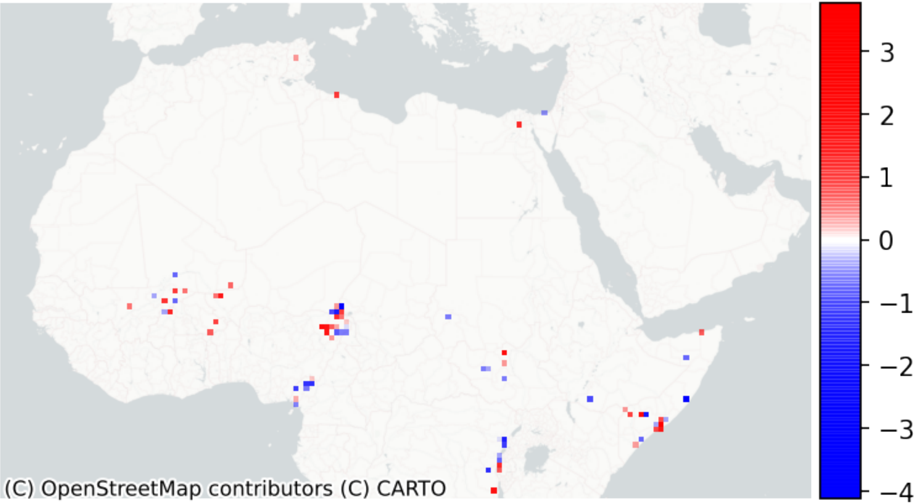

# Predicted: December 2018 (+2 Months)

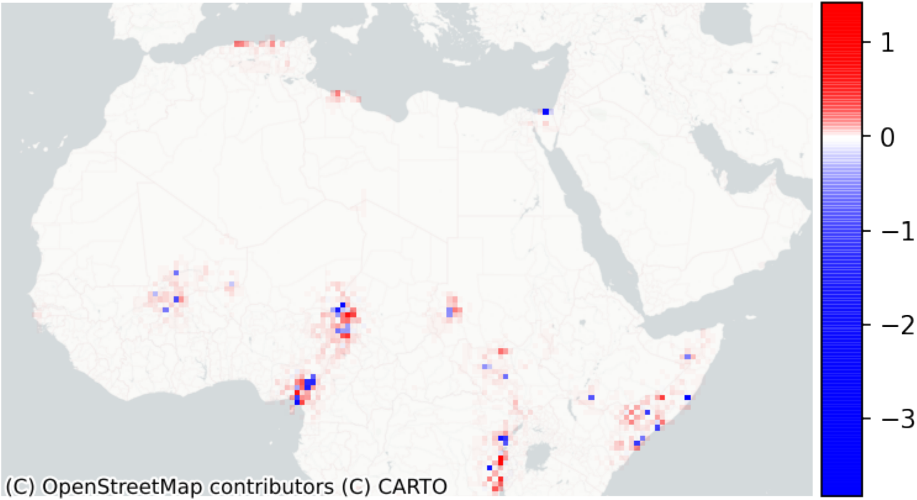

## What's the Model Learning?

**What if...**

- The model is only learning a reversion to the mode (0) when the current death count is greater than 0.
- And, when the current death count is 0, it just predicts something like the mean increase in deaths from the training set?

$$\hat{\Delta}_{s=X} = \begin{cases} -\ln(\text{deaths}+1)_{s=0} & \text{if} \ln(\text{deaths}+1)_{s=0} > 0 \\ \bar{\Delta}_{s \neq X} & \text{else} \end{cases}$$

# Actual versus Predicted (Validation + Test set)

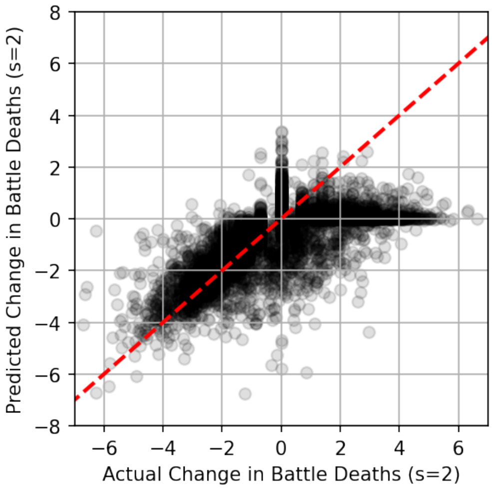


Figure 9: Observed v. Predicted

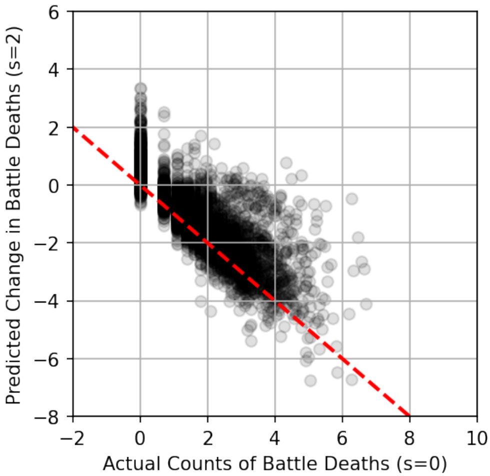


Figure 10: Count v. Prediction

# Black Box

Can we open up the “black box” of the neural network?

## Methods to Inspect Model

- ~~Shapley values~~
- ~~Local Interpretable Model-Agnostic Explanations (LIME)~~
- Attention Layer [Bahdanau et al., 2016]
- Occlusion Sensitivity [Zeiler and Fergus, 2014]
- Alternative Models

# Attention Layer & Occlusion

| Feature | Importance |
|---:|---:|
| Current $\ln(\text{deaths} + 1)$ | 0.284 |
| Population | 0.271 |
| Urban area | 0.207 |
| Travel time to major city | 0.051 |
| Agricultural area | 0.040 |
| Gross cell product | 0.035 |
| Forest area | 0.029 |
| Diamond resources | 0.017 |
| Barren area | 0.016 |
| Border distance | 0.014 |
| Savanna area | 0.012 |
| Pasture area | 0.011 |
| Missingness ind. | 0.010 |
| Distance to capital | 0.010 |

Let's try the same ConvLSTM model with *only* one feature:

$$ln(\text{battle deaths} + 1).$$

| | Competition Model ConvLSTM | | Single Feature ConvLSTM | |
|---|---|---|---|---|
| Steps | MSE | TADDA | MSE | TADDA |
| $s = 2$ | 0.022 | 0.017 | 0.022 | 0.013 |
| $s = 3$ | 0.022 | 0.016 | 0.022 | 0.013 |
| $s = 4$ | 0.022 | 0.016 | 0.022 | 0.014 |
| $s = 5$ | 0.022 | 0.016 | 0.022 | 0.013 |
| $s = 6$ | 0.023 | 0.017 | 0.022 | 0.013 |
| $s = 7$ | 0.023 | 0.017 | 0.022 | 0.014 |


Is this model...bad?

# No. In fact, it is award winning! (Almost)

This ConvLSTM was among the two best models out of eight teams that submitted.

# Radar Plots

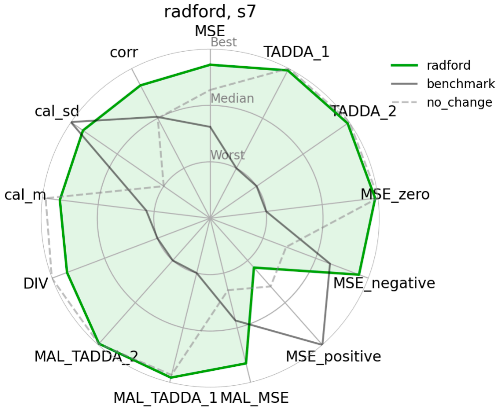


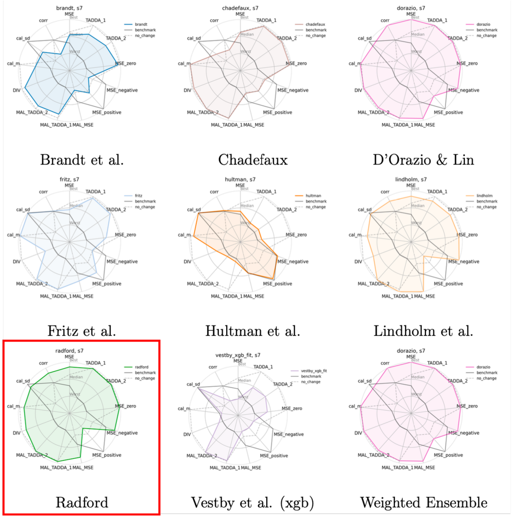


[Paola Vesco and Weidmann, 2022]

# About Those Escalations...

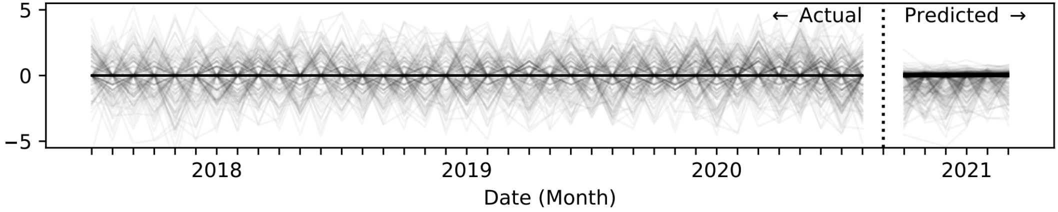


**Figure 11:** Views out-of-sample predictions

Forecasting state-based violence (*escalations*) is really hard.

# Improving Forecasts

- Different models.
- Higher-resolution time-varying predictors.
    - Evidence from Hong Kong Protests.
- Improved measurement.
    - “Estimating Conflict Losses and Reporting Biases”
- Understand the fundamental limits to prediction.

# Models in the Competition

- Elastic net regression
- Random forest
- Graph convolutional neural networks
- Dynamic time warping
- XGboost
- Hierarchical regression models
- State-space models
- Topic models
- AutoML
- Hidden Markov models
- ConvLSTM

- Different models.
- Higher-resolution time-varying predictors.
    - Evidence from Hong Kong Protests.
- Improved measurement.
    - “Estimating Conflict Losses and Reporting Biases”
- Understand the fundamental limits to prediction.

# Hong Kong Protests

# HKMAP.LIVE


Find...
長沙灣
石硤尾
九龍塘
新蒲崗
牛池灣
九龍城
九龍灣
牛頭角
秀茂坪
何文田
土瓜灣
油麻地
觀塘
藍田
九龍灣
紅磡
茶果嶺
北角
鰂魚涌
西灣河
西營盤
石塘咀
上環
堅尼地城
半山區
中環
灣仔
銅鑼灣
9 Oct 18:12
© OpenStreetMap contributors

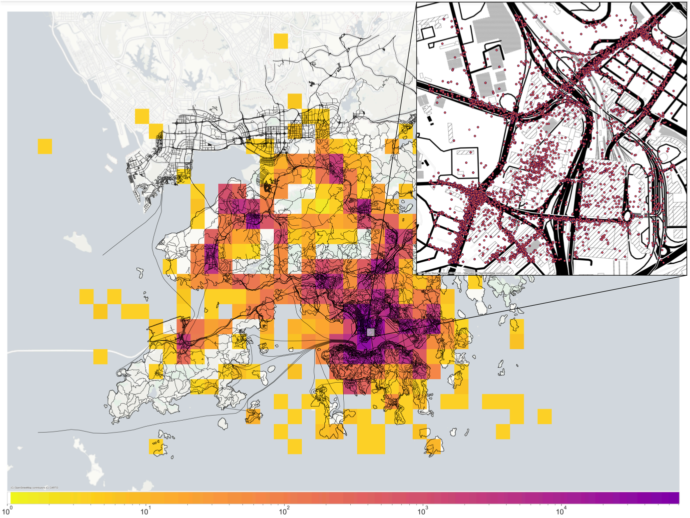
$10^0$
$10^1$
$10^2$
$10^3$
$10^4$

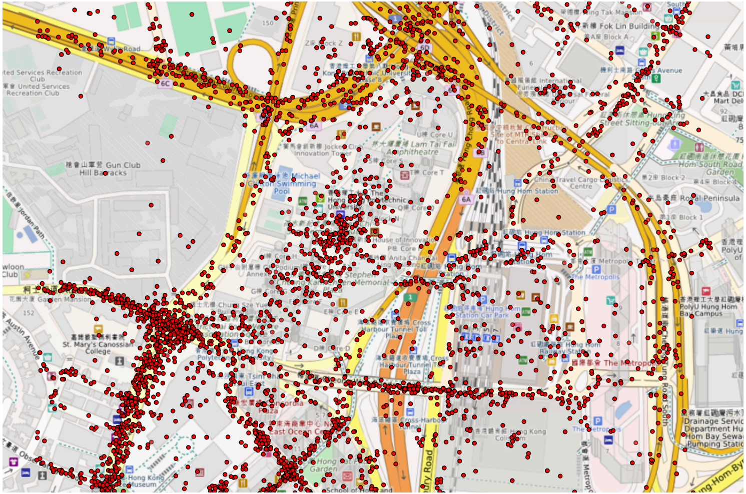

# Hkmap icons

## Approx. 195,262 Reports

| | | | |
|---|---|---|---|
| 🐶 | Police | 🦖 | Special Tactics Unit |
| 👷 | Protesters | 🚫 | Closed |
| 🚓 | Emergency Unit Car | 🚑 | Ambulance |
| ‼️ | Danger | 🚒 | Firetruck |
| 📸 | Live Reporter | ⚠️ | Road Blocked |
| 🏠 | Safe | 💦 | Water Cannon Used |
| 🚩 | Live Bullet Warning | 🟦 | Water Cannon Warning |
| 🏴 | Tear Gas Warning | | |
| 💢 | Tear Gas Released | | |

**Old Question**

How many battle deaths do we expect to observe in this $2500km^2$ cell during one month?

$\downarrow$

**New Question**

What is the probability that at least one of these events is reported within a $4km^2$ grid cell during a 2 hour period?

# Method

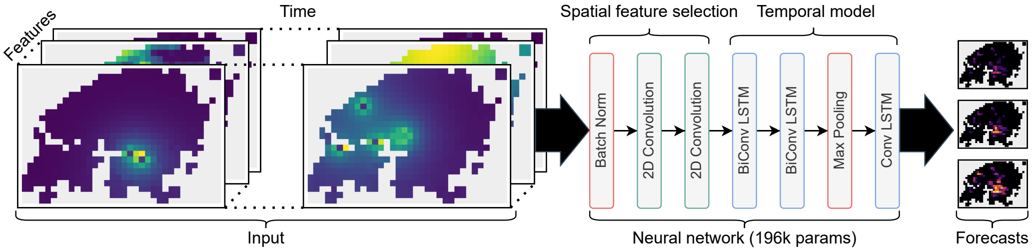

Features
Time
Input
Spatial feature selection
Temporal model
Batch Norm
2D Convolution
2D Convolution
BiConv LSTM
BiConv LSTM
Max Pooling
Conv LSTM
Neural network (196k params)
Forecasts

| Structural Model | Limited Model | Full Model |
|---|---|---|
| Map & Building data | Map & Building data | Map & Building data |
| Roads data | Roads data | Roads data |
| | Weather data | Weather data |
| | Day of week | Day of week |
| | Time of day | Time of day |
| | | HKMap.live reported events |

**Table 1:** Three sets of models.

https://pandas-lab.com/assets/sds_police_2km_2hr.mp4 https://pandas-lab.com/assets/sds_emergency_car_2km_2hr.mp4

# Police Presence P-R Curve

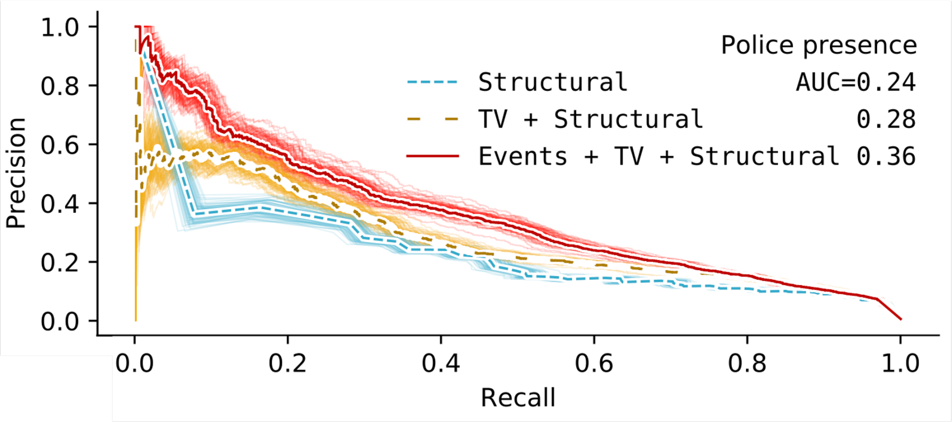

# Swat Team P-R Curve

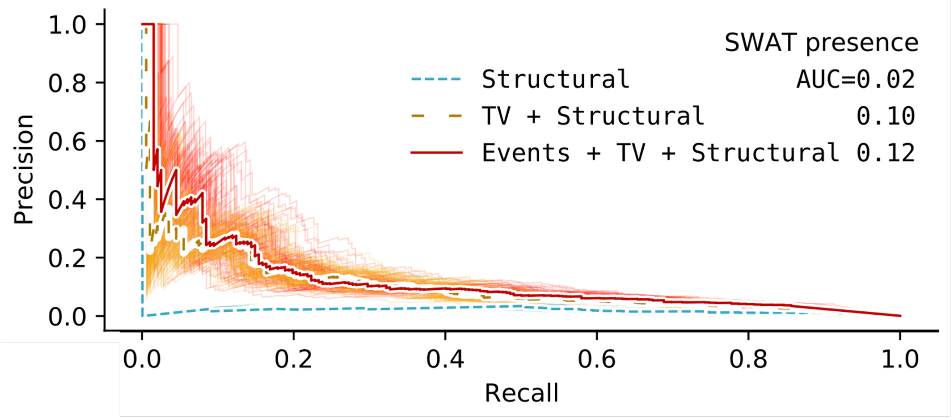

# Time-Varying Predictors Help

## Forecasting HK Events

High-resolution time-varying predictors (e.g., events) help!

## Undercover Police Arrests (with Howard Liu)

- How does the geography of protest shape police tactics?
- We find undercover police are more likely to make arrests on the periphery of protests than in central areas.
- Currently under review...

# Improving Forecasts

- Different models.
- Higher-resolution time-varying predictors.
    - Evidence from Hong Kong Protests.
- Improved measurement.
    - "Estimating Conflict Losses and Reporting Biases"
- Understand the fundamental limits to prediction.

# Measuring Battle Deaths

# UCDP-GED

## UCDP-GED

- Uppsala Conflict Data Program Georeferenced Event Dataset
- Sources:
  - Global newswires & BBC
  - Local and specialized news
  - IGO, NGO, government, and research reports

# Estimating Conflict Losses and Reporting Biases

PNAS

BRIEF REPORT | POLITICAL SCIENCES

## Estimating conflict losses and reporting biases

Benjamin J. Radford[a,b,1], Yaoyao Dai[c], Niklas Stoehr[d], Aaron Schein[e], Mya Fernandez[b,c], and Hanif Sajid[a]

Edited by David Laitin, Stanford University, Stanford, CA; received May 2, 2023; accepted July 17, 2023

Determining the number of casualties and fatalities suffered in militarized conflicts is important for conflict measurement, forecasting, and accountability. However, given the nature of conflict, reliable statistics on casualties are rare. Countries or political actors involved in conflicts have incentives to hide or manipulate these numbers, while third parties might not have access to reliable information. For example, in the ongoing militarized conflict between Russia and Ukraine, estimates of the magnitude of losses vary wildly, sometimes across orders of magnitude. In this paper, we offer an approach for measuring casualties and fatalities given multiple reporting sources and, at the same time, accounting for the biases of those sources. We construct a dataset of 4,609 reports of military and civilian losses by both sides. We then develop a

**Old Question**

How many battle deaths do we expect to observe in this 2500km$^2$ cell during one month?

↓

**New Question**

- Model-based measurement of battle deaths.
- Can we simultaneously account for source biases?

LONDON, Sept 21 (Reuters) - Russia will draft 300,000 reservists to support its military campaign in Ukraine, Defence Minister Sergei Shoigu said on Wednesday in televised remarks.

In Moscow's first update on casualty numbers in almost six months, Shoigu said 5,937 Russian soldiers had been killed since the start of the conflict.

Reuters (Sept. 21, 2022)

# What Does the USG Say?

U.S. Department *of* Defense

His action follows a Ukrainian counteroffensive that pushed Russian forces from Kharkiv and liberated more than 3,000 square kilometers of Ukrainian territory. In August, DOD Policy Chief Colin Kahl said the Russians have lost between 50,000 and 70,000 service members in its war on Ukraine.

U.S. D.O.D. (September 22, 2022)

# What about Ukrainian losses?

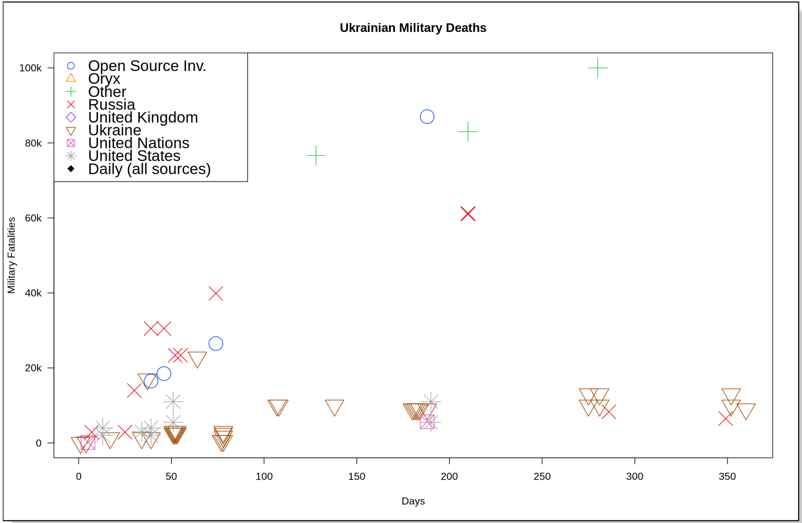

# Method

- Collected 4,609 reports on battle losses (news and social media):
  - Deaths & casualties
  - Aircraft
  - Drones
  - Vehicles
  - Anti-aircraft systems
- Bayesian random effects generalized additive model (Stan)
- Source bias random effects

# The full model

**Likelihood**

$$y_i^{\text{daily}} \sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \quad (1)$$

$$y_j^{\text{cum}} \sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1/\exp(\phi_{ct[j]})) \quad (2)$$

**Loss means**

$$\mu_i^{\text{daily}} = \theta_{ct[i],d[i]} + \beta^{\text{bias}}_{c[i],st[i]} + \beta^{\text{min}}_{s[i]} I_i^{\text{min}} + \beta^{\text{max}}_{s[i]} I_i^{\text{max}} \quad (3)$$

$$\mu_j^{\text{cum}} = \ln(\Sigma_{k=1}^{d[j]} \exp(\theta_{ct[j],d[k]})) + \beta^{\text{bias}}_{c[j],st[j]} + \beta^{\text{min}}_{s[j]} I_j^{\text{min}} + \beta^{\text{max}}_{s[j]} I_j^{\text{max}} \quad (4)$$

**Latent time series**

$$\theta_{ct,d} = (B\beta_{ct}^{\text{spline}})_d + \beta_{ct}^{\text{const}} + \beta_{ct}^{\text{trend}}\,(d/365) \quad (5)$$

**Priors**

$$\beta_c^{\text{const}} \sim N(\mu^{\text{const}}, \sigma^{\text{const}}) \quad (6)$$

$$\beta_{ct}^{\text{trend}} \sim N(\mu^{\text{trend}}, \sigma^{\text{trend}}) \quad (7)$$

$$\beta_{ct}^{\text{spline}} \sim N(0, \Sigma^{\text{spline}}) \quad (8)$$

$$\beta_s^{\text{min}} \sim N(\mu^{\text{min}}, \sigma^{\text{min}}) \quad (9)$$

$$\beta_s^{\text{max}} \sim N(\mu^{\text{max}}, \sigma^{\text{max}}) \quad (10)$$

$$\beta_{c,st}^{\text{bias}} \sim N(\gamma_{st}^{\text{bias}}, \sigma_{st}^{\text{bias}}) \quad (11)$$

$$\gamma_{st}^{\text{bias}} \sim N(0, \sigma_1^{\text{bias}}) \quad (12)$$

$$\phi_{ct} \sim N(\mu^{\phi}, \sigma^{\phi}) \quad (13)$$

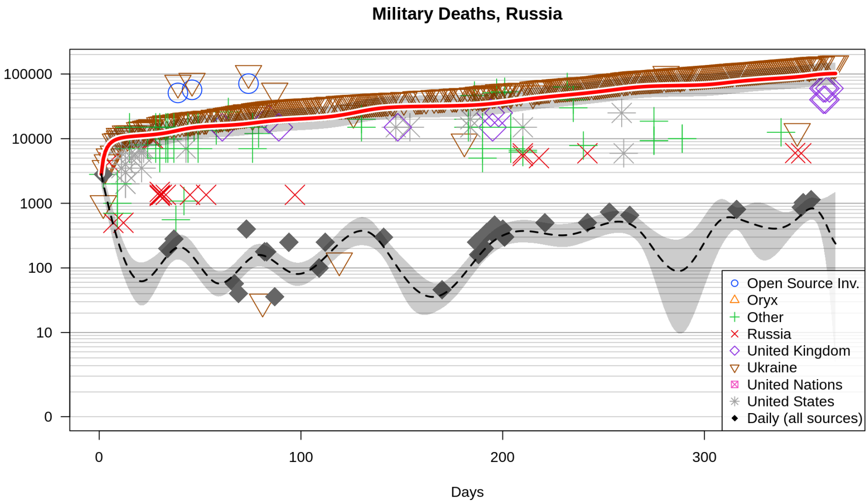
Military Deaths, Russia
100000
10000
1000
100
10
0
0
100
200
300
Days
Open Source Inv.
Oryx
Other
Russia
United Kingdom
Ukraine
United Nations
United States
Daily (all sources)

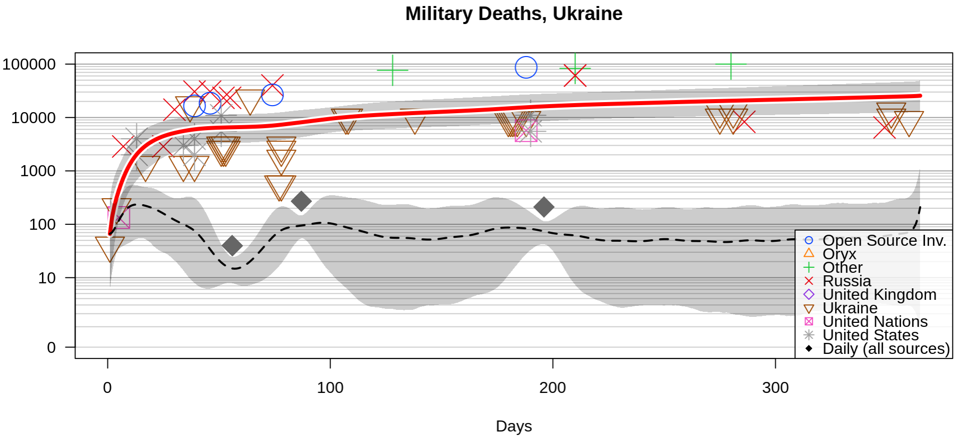

Military Deaths, Ukraine
100000
10000
1000
100
10
0
0
100
200
300
Days
Open Source Inv.
Oryx
Other
Russia
United Kingdom
Ukraine
United Nations
United States
Daily (all sources)

- Expected deaths as of day 365

| | | | | |
|---|---|---|---|---|
| RU | 76,687 | (38,670 | – | 139,772) |
| UA | 17,223 | (6,219 | – | 39,105) |

- Expected deaths as of day 365

| | | | | |
|---|---|---|---|---|
| RU | 76,687 | (38,670 | – | 139,772) |
| UA | 17,223 | (6,219 | – | 39,105) |

- Casualties to deaths ratios on day 365

| | |
|---|---|
| RU | 2.9:1 |
| UA | 4.9:1 |

## Interesting numbers

- Expected deaths as of day 365

| | | | | |
|---|---|---|---|---|
| RU | 76,687 | (38,670 | – | 139,772) |
| UA | 17,223 | (6,219 | – | 39,105) |

- Casualties to deaths ratios on day 365

| | |
|---|---|
| RU | 2.9:1 |
| UA | 4.9:1 |

- Russian to Ukrainian troop loss ratio

  5.53:1 (1.61:1 – 14.5:1)

- Expected deaths as of day 365

| | | | | |
|---|---|---|---|---|
| RU | 76,687 | (38,670 | – | 139,772) |
| UA | 17,223 | (6,219 | – | 39,105) |

- Casualties to deaths ratios on day 365

| | |
|---|---|
| RU | 2.9:1 |
| UA | 4.9:1 |

- Russian to Ukrainian troop loss ratio

  5.53:1 (1.61:1 – 14.5:1)

# Biases

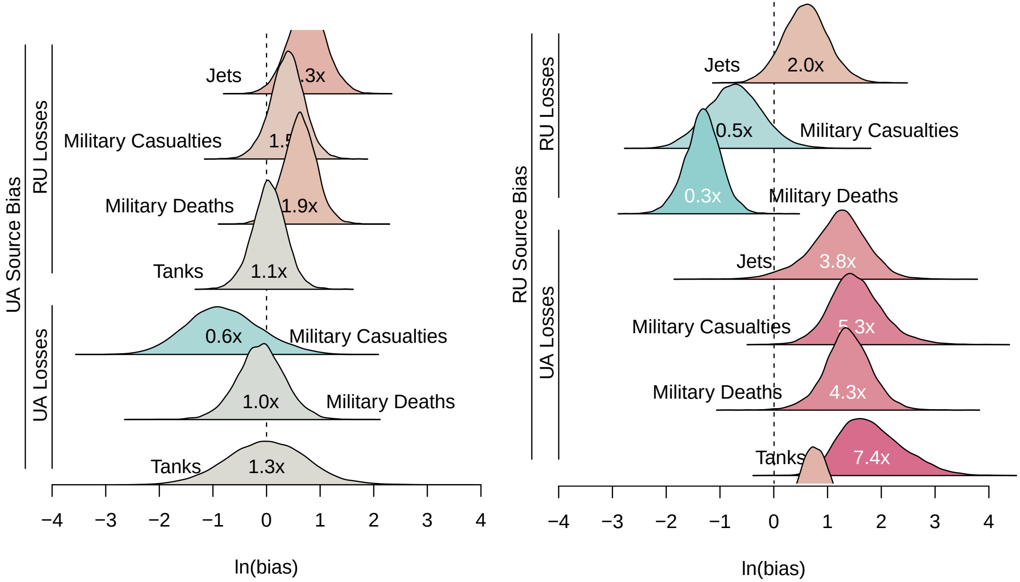

UA Source Bias
RU Losses
Jets
.3x
Military Casualties
1.5
Military Deaths
1.9x
Tanks
1.1x
UA Losses
0.6x
Military Casualties
1.0x
Military Deaths
Tanks
1.3x
−4
−3
−2
−1
0
1
2
3
4
ln(bias)
RU Source Bias
RU Losses
Jets
2.0x
0.5x
Military Casualties
0.3x
Military Deaths
UA Losses
Jets
3.8x
Military Casualties
5.3x
Military Deaths
4.3x
Tanks
7.4x
−4
−3
−2
−1
0
1
2
3
4
ln(bias)

# Improving Forecasts

- Different models.
- Higher-resolution time-varying predictors.
    - Evidence from Hong Kong Protests.
- Improved measurement.
    - "Estimating Conflict Losses and Reporting Biases"
- Understand the fundamental limits to prediction.

# Conclusion

# Conclusion

## The Message

- Forecasting is important for science!
  - Falsify theories.
  - Predict effects of interventions.
  - Avoid overfitting.
- Forecasting *conflict escalation* is **really hard**.

## The Good News

- There's never been a better time to study forecasting & measurement.
- Interesting data are more available than ever.
- Social science is benefiting machine learning and data science best practices.

# Additional Projects

## Measurement

- Estimating complex concepts from text data:
  - Populism
  - Political ideology
  - Event attributes
- Measuring sub-national territorial control

## Data Science for Cybersecurity

- Using language models to identify vulnerabilities in source code
- Machine learning to detect network intrusions

# Thank You

Benjamin Radford
benjamin.radford@charlotte.edu

D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate, 2016. URL `https://arxiv.org/abs/1409.0473`.

F. Chollet. On the measure of intelligence, 2019. URL `https://arxiv.org/abs/1911.01547`.

M. Colaresi and Z. Mahmood. Do the robot: Lessons from machine learning to improve conflict forecasting. Journal of Peace Research, 54(2):193–214, 2017.

E. Gartzke. War Is in the Error Term. International Organization, 53(3):567–587, 1999.

H. Hegre, P. Vesco, and M. Colaresi. Lessons from an escalation prediction competition. International Interactions, 48(4):521–554, 2022.

J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models, 2020. URL `https://arxiv.org/abs/2001.08361`.

W. Lotter, G. Kreiman, and D. Cox. Deep predictive coding networks for video prediction and unsupervised learning, 2017. URL `https://arxiv.org/abs/1605.08104`.

D. Muchlinski, D. Siroky, J. He, and M. Kocher. Comparing random forest with logistic regression for predicting class-imbalanced civil war onset data. Political Analysis, 24(1):87–103, 2016. doi: 10.1093/pan/mpv024.

M. C. R. B. J. A. L. G. R. Paola Vesco, Håvard Hegre and N. B. Weidmann. United they stand: Findings from an escalation prediction competition. International Interactions, 48(4):860–896, 2022.

R. Sundberg and E. Melander. Introducing the UCDP Georegerenced Event Dataset. Journal of Peace Research, 50 (4):523–532, 2013.

A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In T. Linzen, G. Chrupała, and A. Alishahi, editors, Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP, pages 353–355, Brussels, Belgium, Nov. 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5446. URL https://aclanthology.org/W18-5446.

M. D. Ward, N. W. Metternich, C. L. Dorff, M. Gallop, F. M. Hollenbach, A. Schultz, and S. Weschle. Learning from the Past and Stepping into the Future: Toward a New Generation of Conflict Prediction. International Studies Review, 15(4):473–490, 12 2013. ISSN 1521-9488. doi: 10.1111/misr.12072. URL https://doi.org/10.1111/misr.12072.

M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, Computer Vision – ECCV 2014, pages 818–833, Cham, 2014. Springer International Publishing.

# International Interactions

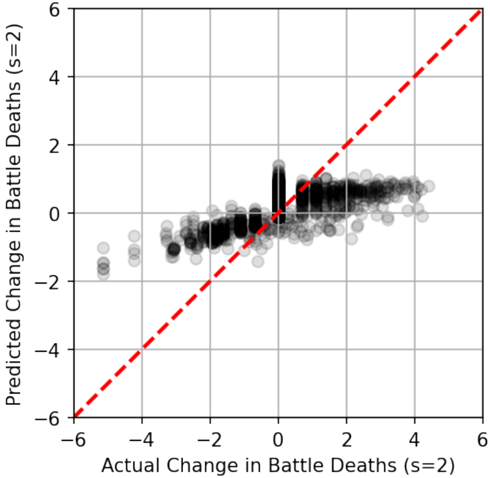


Figure 12: Benchmark Model

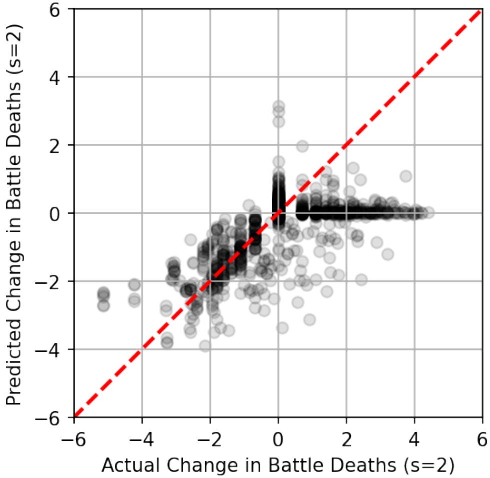


Figure 13: ConvLSTM

Out-of-sample forecasts

# The Numbers

**Table 2:** Mean Sq. Error

| Steps | ConvLSTM | Benchmark |
|---|---|---|
| $s = 2$ | 0.024 | 0.053 |
| $s = 3$ | 0.028 | 0.059 |
| $s = 4$ | 0.025 | 0.052 |
| $s = 5$ | 0.036 | 0.064 |
| $s = 6$ | 0.036 | 0.063 |
| $s = 7$ | 0.035 | 0.064 |

**Table 3:** TADDA

| Steps | ConvLSTM | Benchmark |
|---|---|---|
| $s = 2$ | 0.018 | 0.138 |
| $s = 3$ | 0.019 | 0.142 |
| $s = 4$ | 0.020 | 0.142 |
| $s = 5$ | 0.024 | 0.150 |
| $s = 6$ | 0.021 | 0.148 |
| $s = 7$ | 0.025 | 0.153 |

*Targeted Abs. Distance with Direction Augmentation

# Motivation

- Large language models are expensive.
- Companies want to predict their return on performance.
- Researchers have found limits to their ability to predict language [Kaplan et al., 2020].
  - Irreducible error
  - Bayes error

**We can apply this to social sciences:**

Holding data size, compute resources, or model parameters constant, **estimate the limit** of our **ability to make predictions**.

# Hong Kong

# Train / Eval / Test

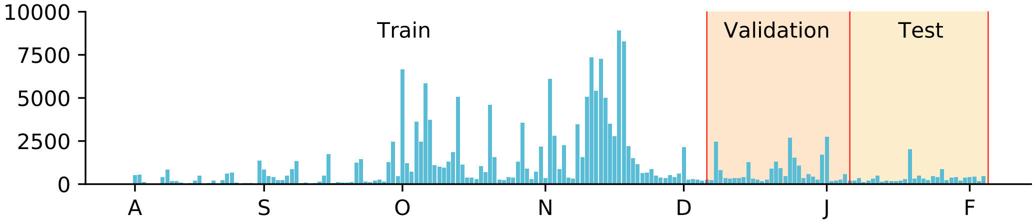


**Figure 14:** HKMap.live data over time

# Police Vehicles P-R Curve

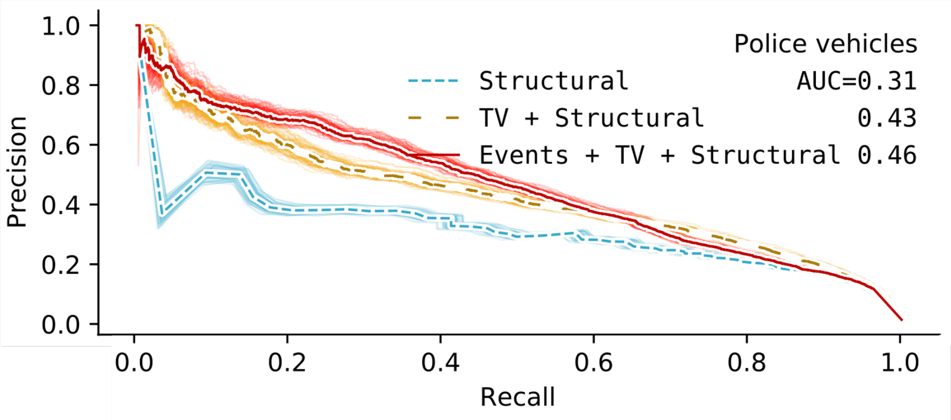

# Tear Gas P-R Curve

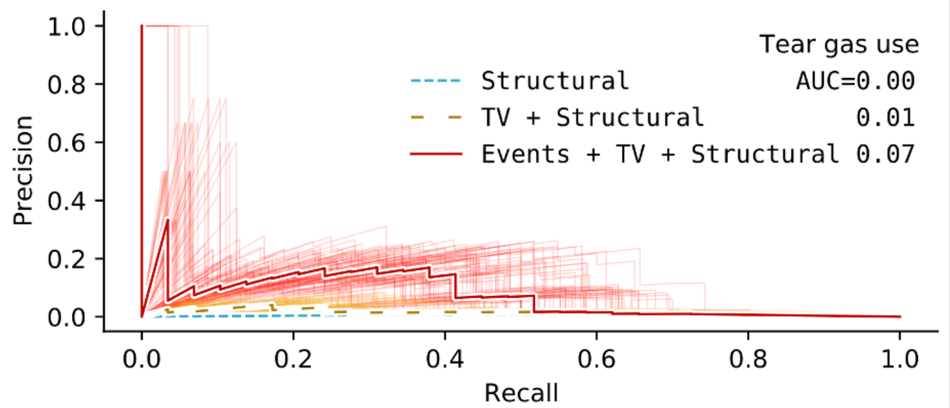

# Do Time-Varying Predictors Help?

| | Police Presence | |
|---|---|---|
| | Limited Model | Full Model |
| Accuracy | 0.92 | **0.93** |
| Precision | 0.07 | 0.07 |
| Recall | **0.95** | 0.91 |
| F-score | 0.13 | **0.14** |

n = 1766 / 289,138

**Table 4:** Evaluation set performance.

# Do Time-Varying Predictors Help?

| | Water Cannon Use | |
|---|---|---|
| | Limited Model | Full Model |
| Accuracy | **0.98** | 0.97 |
| Precision | **0.02** | 0.01 |
| Recall | 0.57 | **0.71** |
| F-score | 0.02 | 0.02 |

n = 97 / 289,138

**Table 5:** Evaluation set performance.

PNAS

# Sources aren't even internally consistent

FORBES > BUSINESS

BREAKING

## 500 Or 10,000 Deaths? Russian Media Finally Seems To Report Dire Troop Casualty Numbers—And Then Deletes Them

**Mason Bissada** Former Staff

Forbes (March 22, 2022)

# Even more data

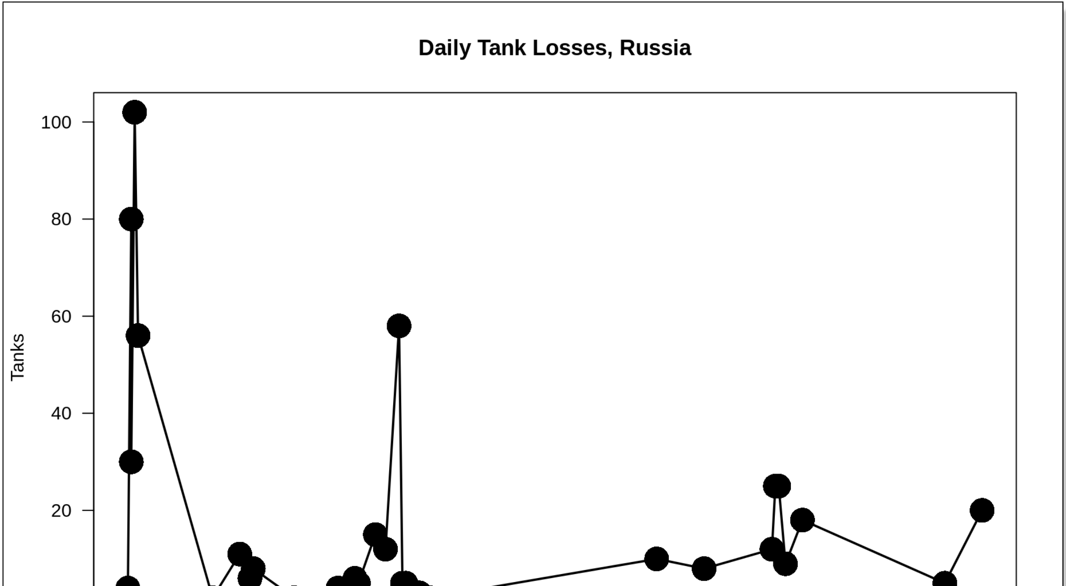

# Putting it all together

Can we put all of this messy data together to obtain nice (“unbiased”) estimates of losses?

# Data

We collected **4,357*** loss reports from Feb. 24, 2022 – Feb. 23, 2023.

- Date
- Loss type (21 types)
- Temporal unit (cumulative or daily)
- Ranges ("between XXX and YYY losses")
- Reporting sources
- Loss country (Ukraine or Russia)
- Newspaper or venue
- Text of claim

* 6,032

# The model diagram

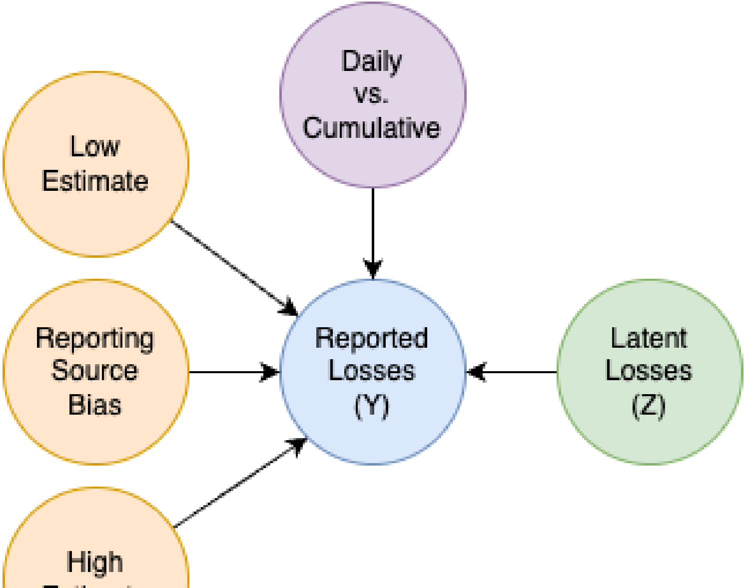

# The model

- Bayesian random effects model (Stan)
- Outcomes:
  - Daily counts $\sim$ Poisson
  - Cumulative counts $\sim$ Negative binomial
- Predictors:
  - Daily latent time series for every type-target pair
  - Reporting source-target biases
  - Reporting source specific min and max scalar

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\ \vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\quad \vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta^{\text{min}}_{\text{source}[i]} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta^{\text{max}}_{\text{source}[i]} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta^{\text{min}}_{\text{source}[i]} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \color{red}{\beta^{\text{max}}_{\text{source}[i]} \times 1_{\text{max}}(x_i^{\text{max}})} \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# The model

$$
\begin{aligned}
y_i^{\text{daily}} &\sim \text{Pois}(\exp(\mu_i^{\text{daily}})) \\
y_j^{\text{cum}} &\sim \text{NB}(\exp(\mu_j^{\text{cum}}), 1\exp(\phi_{\text{type-target}[j]})) \\
\mu_i^{\text{day}} &= \theta_{\text{type-target[i]}} \\
&\quad + \beta_{\text{source-target}[i]} \\
&\quad + \beta_{\text{source}[i]}^{\text{min}} \times 1_{\text{min}}(x_i^{\text{min}}) \\
&\quad + \beta_{\text{source}[i]}^{\text{max}} \times 1_{\text{max}}(x_i^{\text{max}}) \\
\mu_j^{\text{cum}} &= \ln(\text{CumSum}(\exp(\theta_{\text{type-target[j]}}))) \\
&\vdots
\end{aligned}
$$

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate
- Reports are sometimes given as ranges
  - Dummy variables for min or max estimates

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate
- Reports are sometimes given as ranges
  - Dummy variables for min or max estimates
- Reporting sources are likely biased
  - Multilevel bias terms
  - Every reporting source gets a bias term
  - Every loss type within a source gets its own bias term

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate
- Reports are sometimes given as ranges
  - Dummy variables for min or max estimates
- Reporting sources are likely biased
  - Multilevel bias terms
  - Every reporting source gets a bias term
  - Every loss type within a source gets its own bias term
- Biases are multipliers, not intercept shifts
  - Our model is log-linear in parameters

# Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate
- Reports are sometimes given as ranges
  - Dummy variables for min or max estimates
- Reporting sources are likely biased
  - Multilevel bias terms
  - Every reporting source gets a bias term
  - Every loss type within a source gets its own bias term
- Biases are multipliers, not intercept shifts
  - Our model is log-linear in parameters
- Missing values
  - Implicit imputation

## Our special special model

- Multivariate outcomes
  - Daily outcomes
  - Cumulative outcomes
- Actual losses are latent
  - Cubic B-splines to interpolate
- Reports are sometimes given as ranges
  - Dummy variables for min or max estimates
- Reporting sources are likely biased
  - Multilevel bias terms
  - Every reporting source gets a bias term
  - Every loss type within a source gets its own bias term
- Biases are multipliers, not intercept shifts
  - Our model is log-linear in parameters
- Missing values
  - Implicit imputation

Estimating this model was a nightmare.

# Results

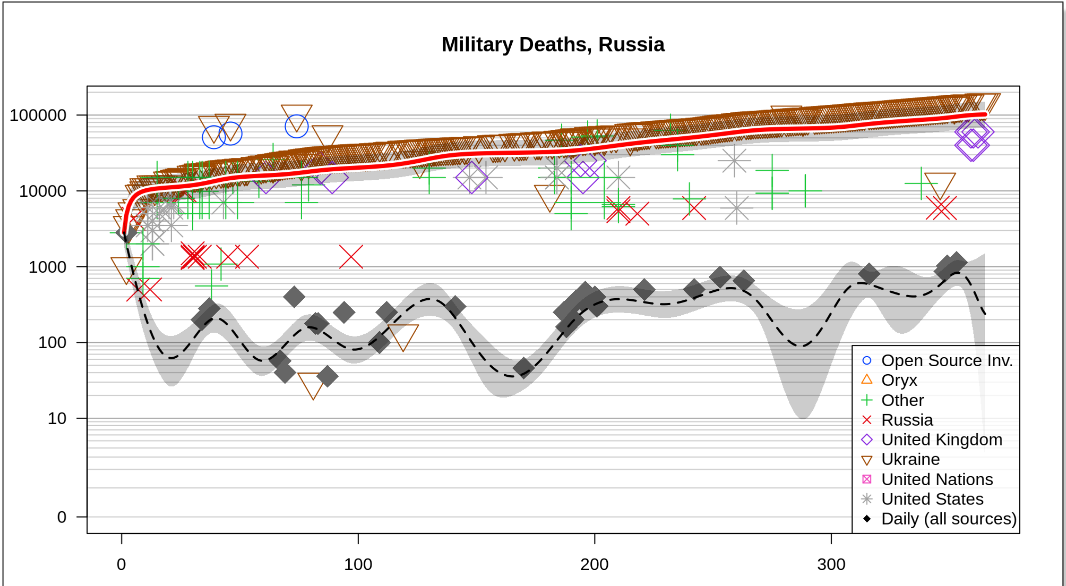
Military Deaths, Russia
100000
10000
1000
100
10
0
0
100
200
300
Open Source Inv.
Oryx
Other
Russia
United Kingdom
Ukraine
United Nations
United States
Daily (all sources)

# Posterior predictive intervals

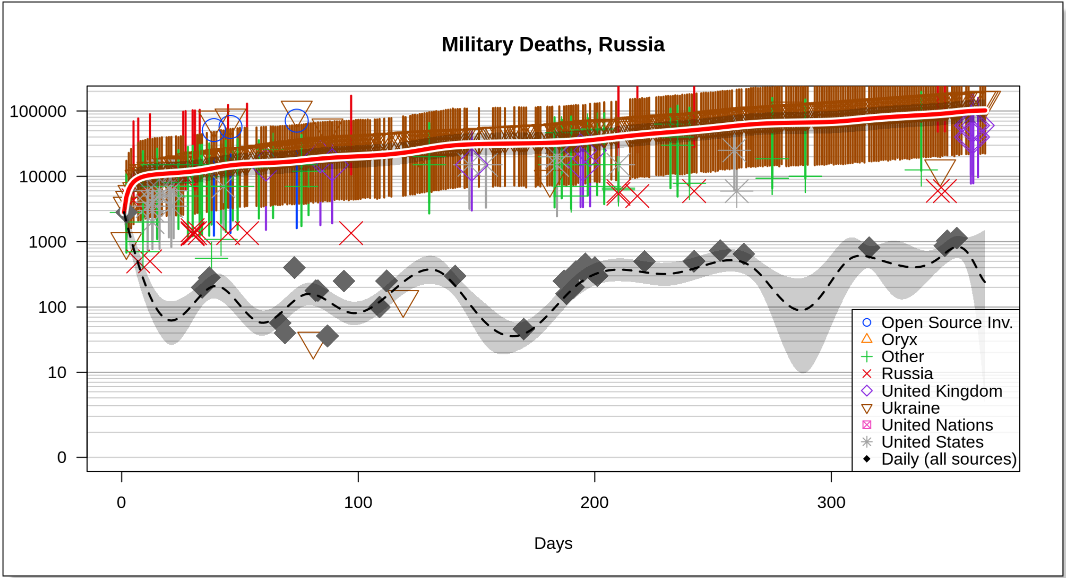

# Ukrainian military fatalities

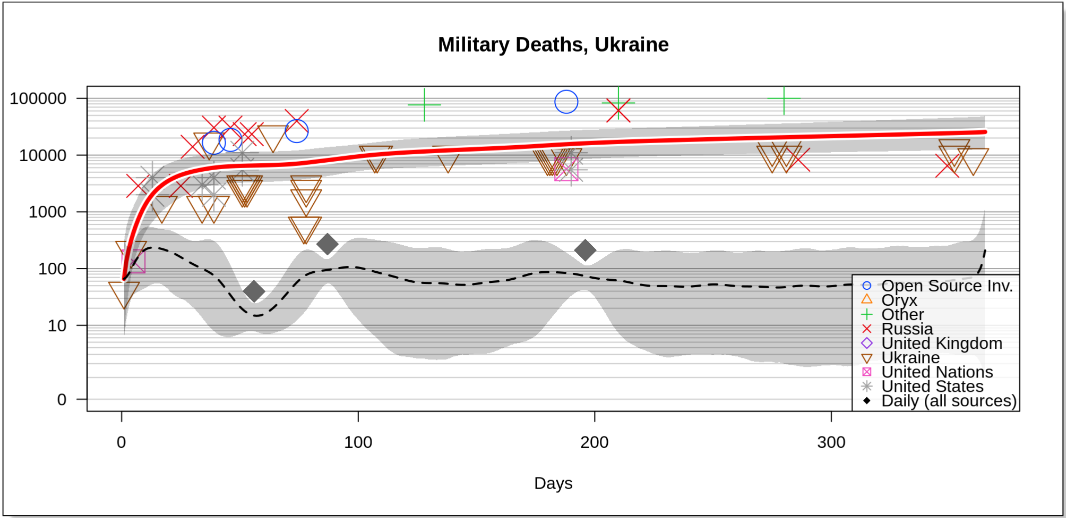

# Tank losses

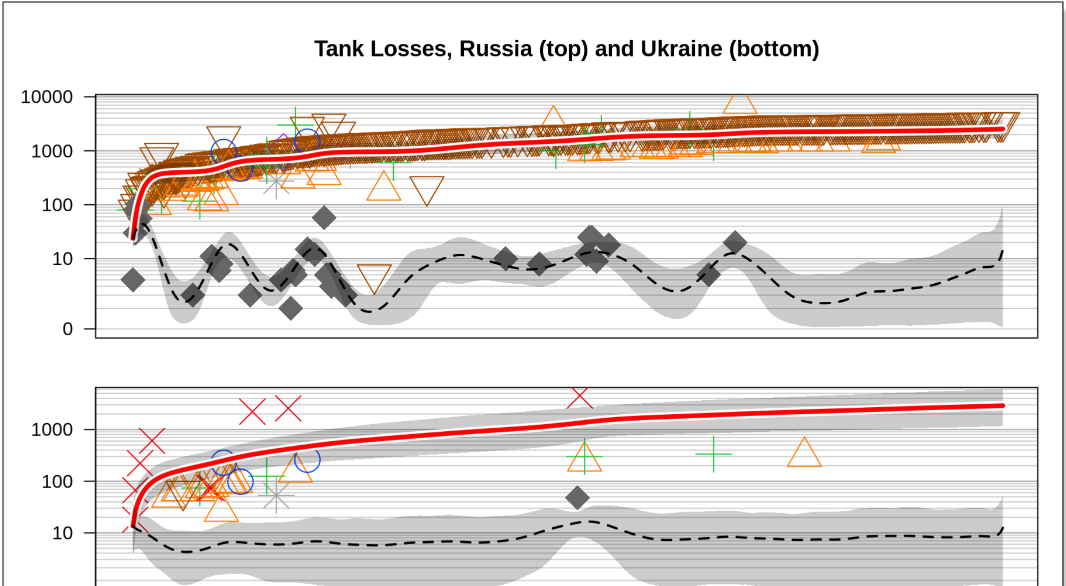

Tank Losses, Russia (top) and Ukraine (bottom)
10000
1000
100
10
0
1000
100
10

# Lots of types

| ISO2 | Loss type | *n* | Est. | 95% CI |
|---|---|---|---|---|
| russian-gray RU | AA Systems | 233 | 339 | [76–1070] |
| ukrainian-gray UA | AA Systems | 13 | 1105 | [108–5247] |
| russian-gray RU | Artillery | 380 | 1483 | [701–2818] |
| ukrainian-gray UA | Artillery | 35 | 2290 | [519–6966] |
| russian-gray RU | Helicopters | 389 | 172 | [87–311] |
| ukrainian-gray UA | Helicopters | 30 | 64 | [14–183] |
| russian-gray RU | Jets | 409 | 146 | [68–273] |
| ukrainian-gray UA | Jets | 38 | 122 | [32–372] |
| russian-gray RU | Military Deaths | 523 | 76687 | [38670–139772] |
| ukrainian-gray UA | Military Deaths | 67 | 17223 | [6219–39105] |
| russian-gray RU | MLRS | 261 | 488 | [148–1222] |
| ukrainian-gray UA | MLRS | 27 | 538 | [155–1482] |
| russian-gray RU | Tanks | 501 | 3380 | [1704–6178] |
| ukrainian-gray UA | Tanks | 33 | 2051 | [385–5946] |
| russian-gray RU | UAVs | 292 | 337 | [153–707] |
| ukrainian-gray UA | UAVs | [illegible] | [illegible] | [illegible] |

# Model validation

- We don't know the true latent values, so we can't evaluate that way.
    - ...or do we?
- There is not a straight-forward way to compute $R^2$ with these models.
- We can compute AIC and BIC, but those only help internally.
- Let's use out-of-sample cross-validation!
    - Divide the data into 5 folds.
    - Estimate 5 models, leaving out one fold of the data each time.
    - Make predictions for $y$ on the held out data.
    - Plot the out-of-sample predictions versus the actual values.

# Model validation

5-fold Cross Validation

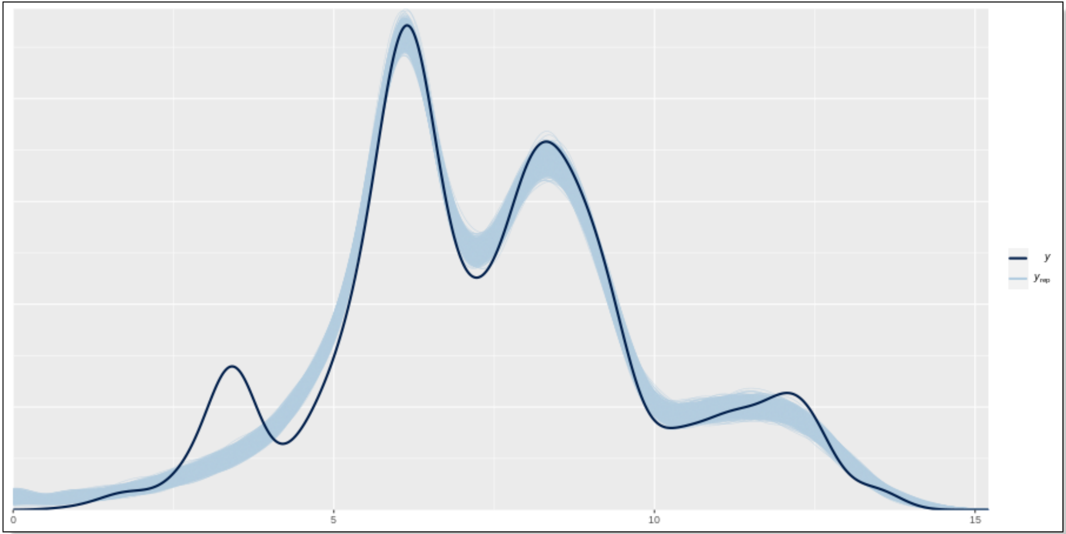

# Model validation

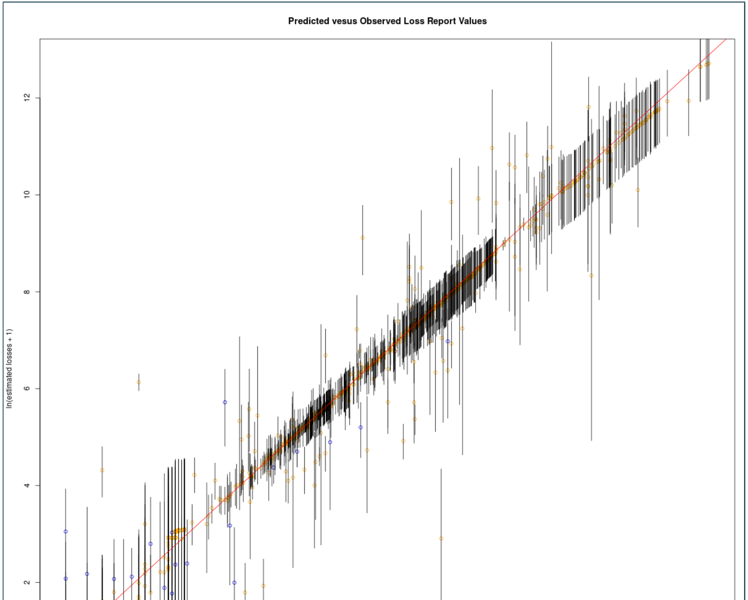

# The Tools

- R / RStudio
  - Data cleaning
  - Graphics and tables
- Stan
  - Bayesian modeling
- University Research Computing
- $\LaTeX$
  - Overleaf.com
  - Manuscript
  - Slides
  - Poster
- Git / Github
  - Collaboration
  - Replication archive
- Harvard Dataverse
  - Data distribution

# Published

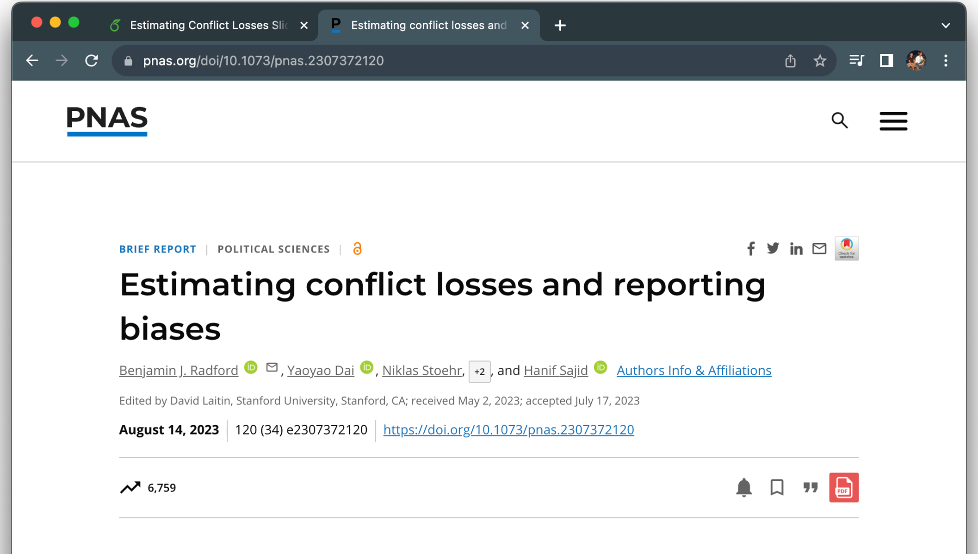
Estimating Conflict Losses Sli
Estimating conflict losses and
pnas.org/doi/10.1073/pnas.2307372120

PNAS

BRIEF REPORT | POLITICAL SCIENCES

# Estimating conflict losses and reporting biases

Benjamin J. Radford, Yaoyao Dai, Niklas Stoehr, +2, and Hanif Sajid Authors Info & Affiliations

Edited by David Laitin, Stanford University, Stanford, CA; received May 2, 2023; accepted July 17, 2023

**August 14, 2023** | 120 (34) e2307372120 | https://doi.org/10.1073/pnas.2307372120

# Timeline

| | |
|---|---|
| **February 24, 2022**: | Started data collection |
| **Fall, 2022**: | Started modeling |
| **February 23, 2023**: | Ended data collection* |
| **May 2, 2023**: | First submission to PNAS |
| **June 5, 2023**: | First round of reviews |
| | 24 page response letter |
| **July 3, 2023**: | Second round of reviews |
| | 18 page response letter |
| **July 20, 2023**: | Acceptance |

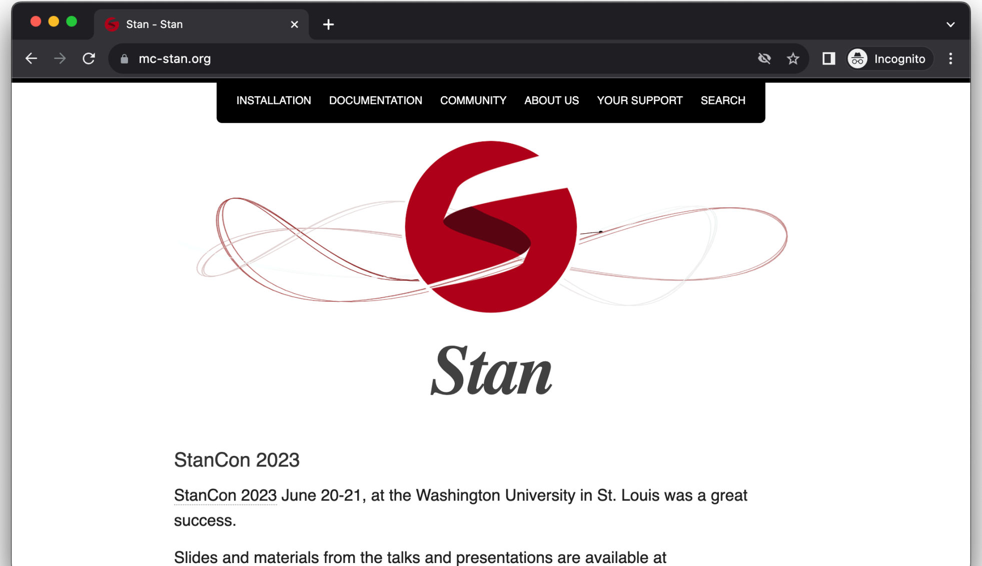
Stan - Stan
mc-stan.org
Incognito
INSTALLATION
DOCUMENTATION
COMMUNITY
ABOUT US
YOUR SUPPORT
SEARCH
Stan
StanCon 2023
StanCon 2023 June 20-21, at the Washington University in St. Louis was a great success.
Slides and materials from the talks and presentations are available at

# Publicity


Wizard_of_Armageddon @Wizard_of_A · Oct 6

Unanimous opinion: this article is garbage. Totally GIGO(Garbage In, Garbage Out).

Wizard_of_Armageddon @Wizard_of_A · Oct 5

The editor-in-chief asked me to review this article (because of my background in statistics). Currently viewing.
theloop.ecpr.eu/estimating-tro...

2 6 12 4,037

# The model, again

- The model is just too complicated
  - We could have made it simpler... probably.
- It's not efficient:
  - could be the Negative Binomial / Poisson link functions
  - could be the $\exp(\ldots)$ operator
- Should we assume:
  - when data are scarce, losses regress to zero
  - when data are scarce, losses regress to their mean

# Please help!

- Is there interest in this in political science?
    - If so, where?
    - Any suggestions for framing?
- Is there interest in follow-up on:
    - estimating daily “conflict intensity”
    - estimating source reporting biases
- Is there anything else we could use our data for?
- Do you think there’s interest in a continuing dataset?

# Conclusion

**Thank you!**

**And thank you to my co-authors:**

Yaoyao Dai

Niklas Stoehr

Aaron Schein

Mya Fernandez

Hanif Sajid

**Contact me:** bradfor7@uncc.edu

# Appendix

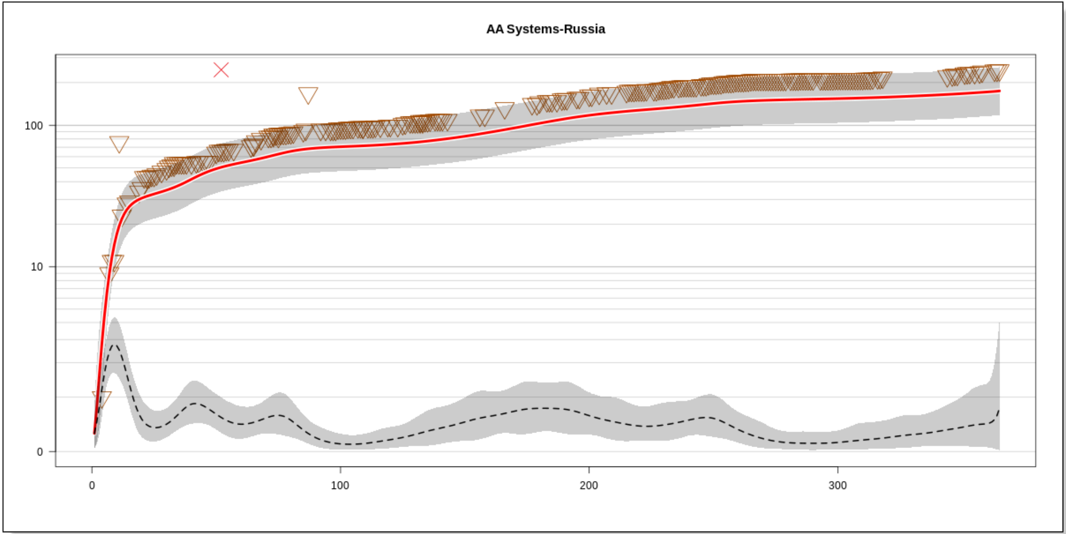

# Appendix

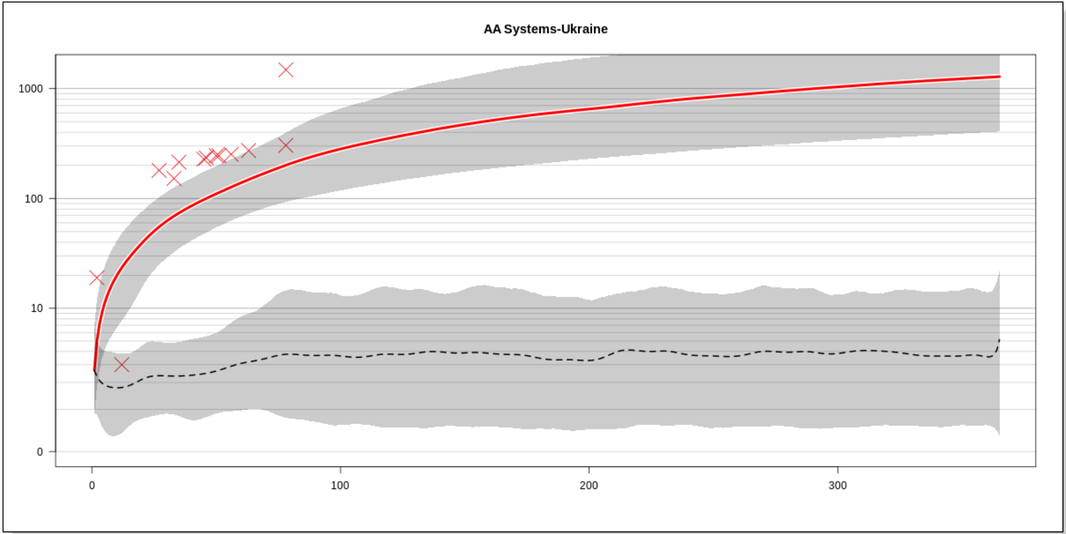

# Appendix

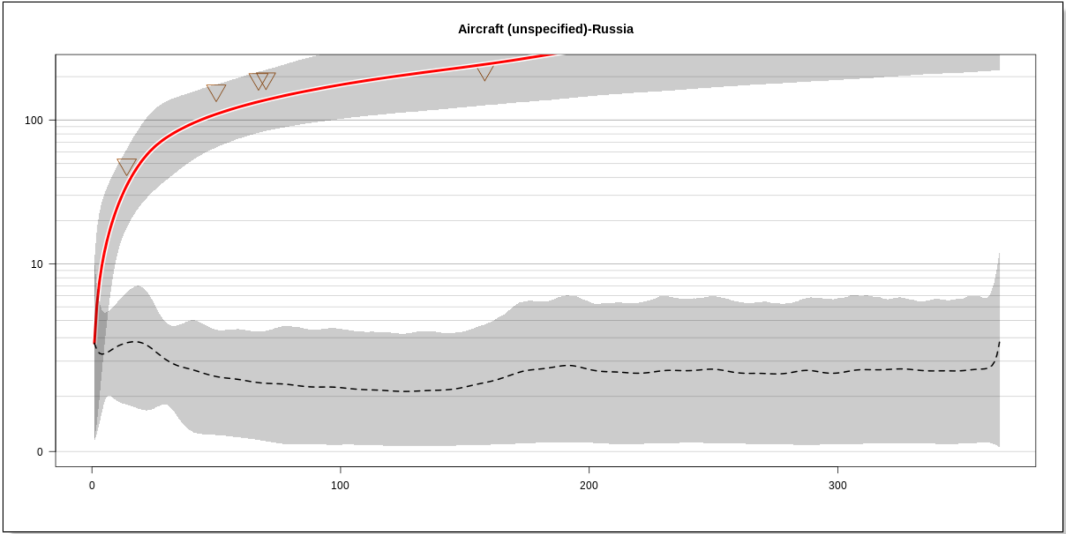

# Appendix

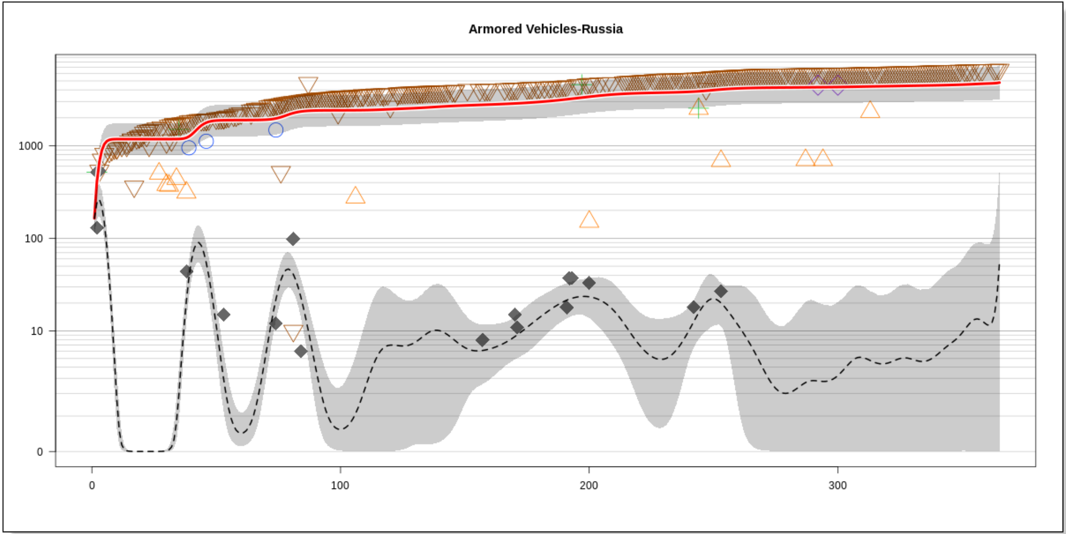

# Appendix

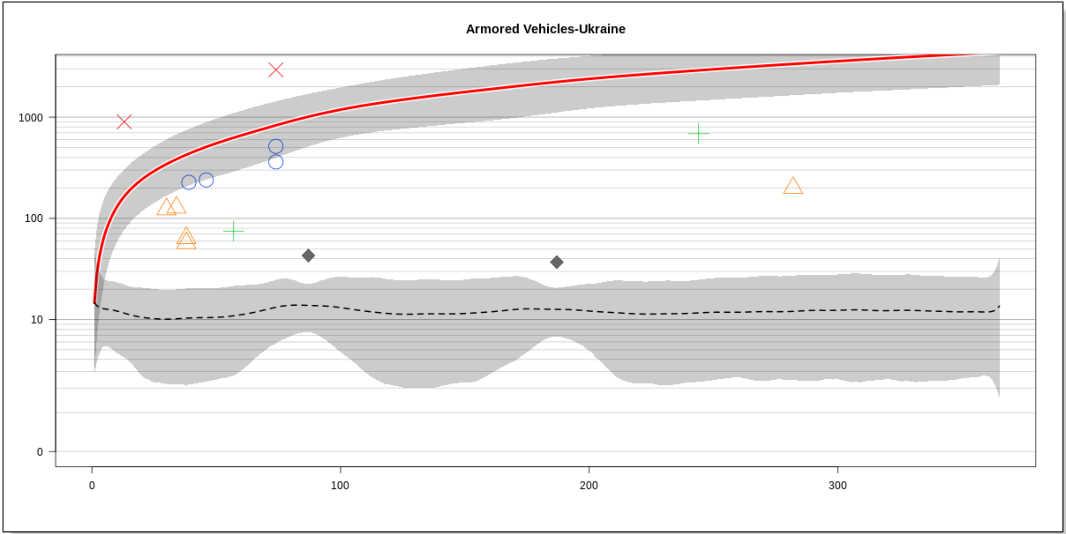

# Appendix

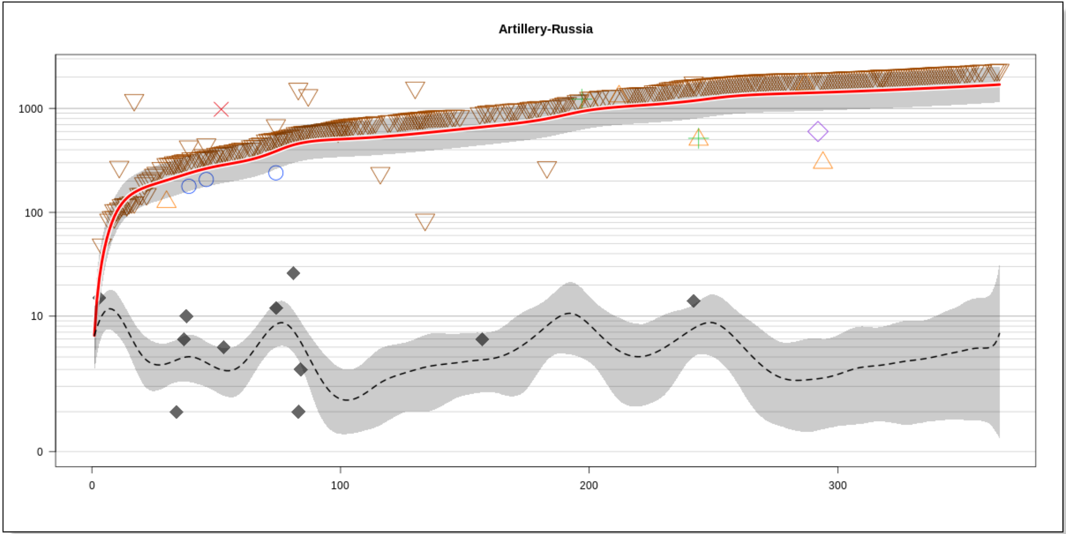

# Appendix

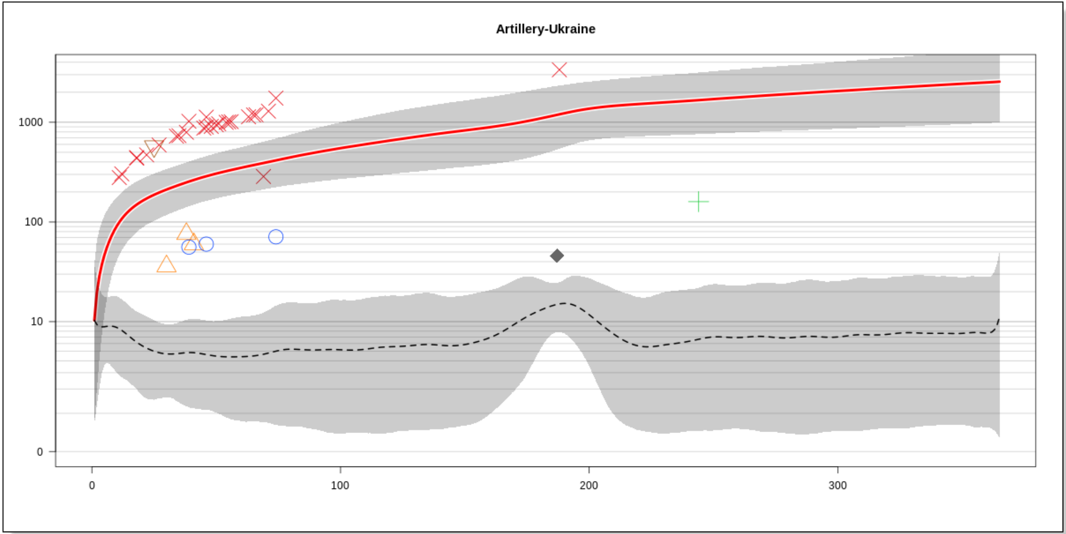

# Appendix

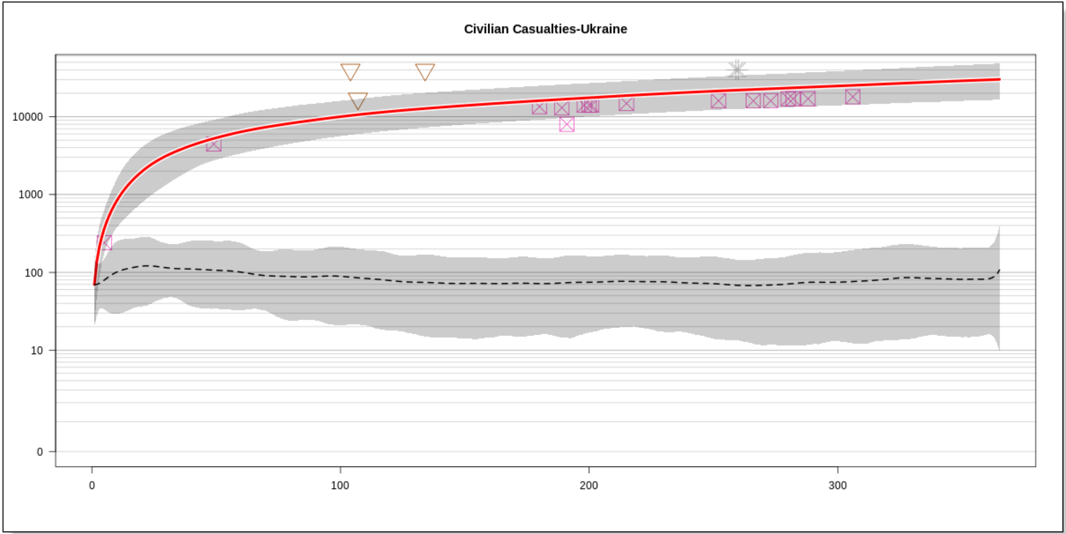

Civilian Casualties-Ukraine
10000
1000
100
10
0
0
100
200
300

# Appendix

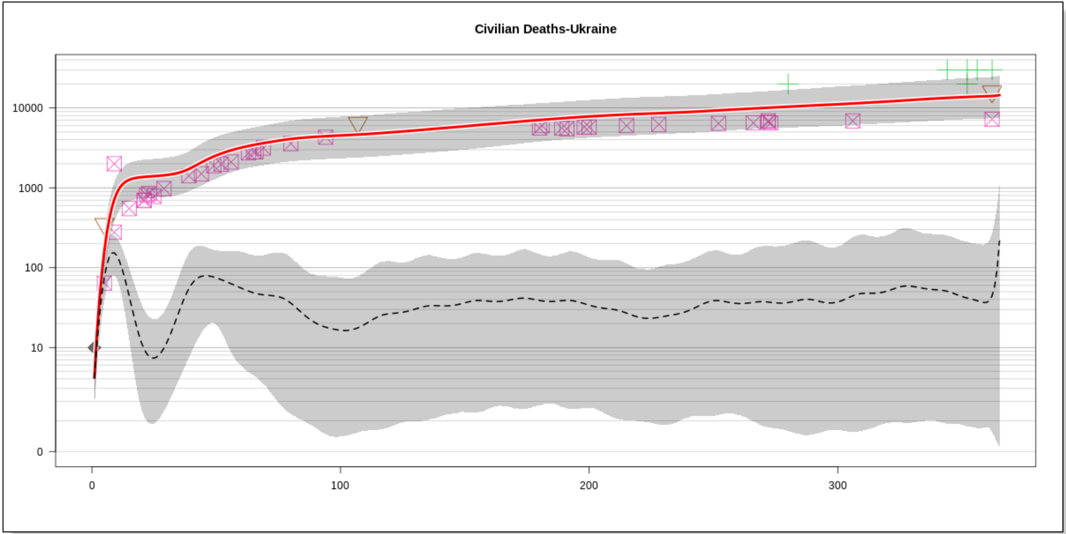

# Appendix

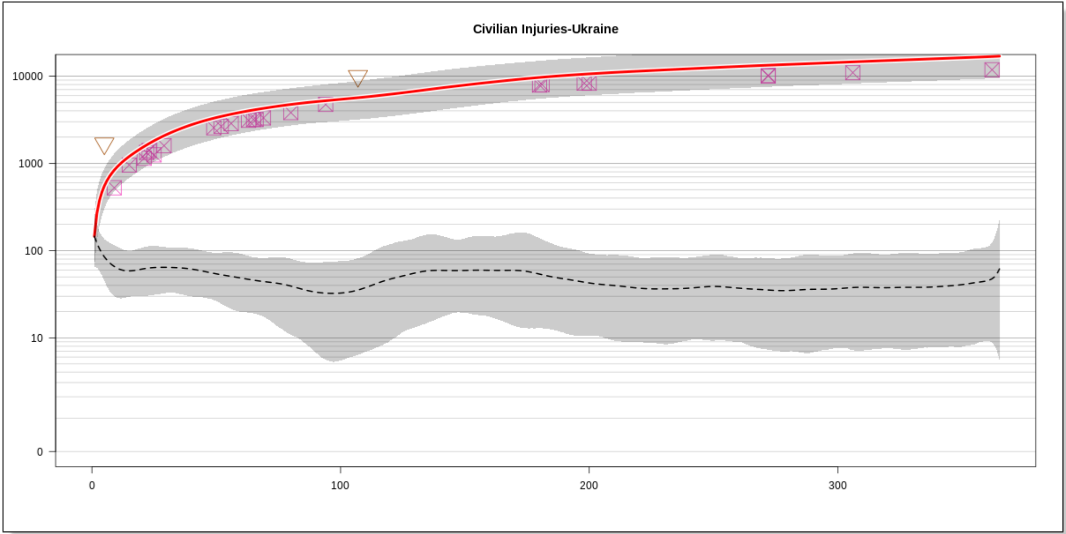

# Appendix

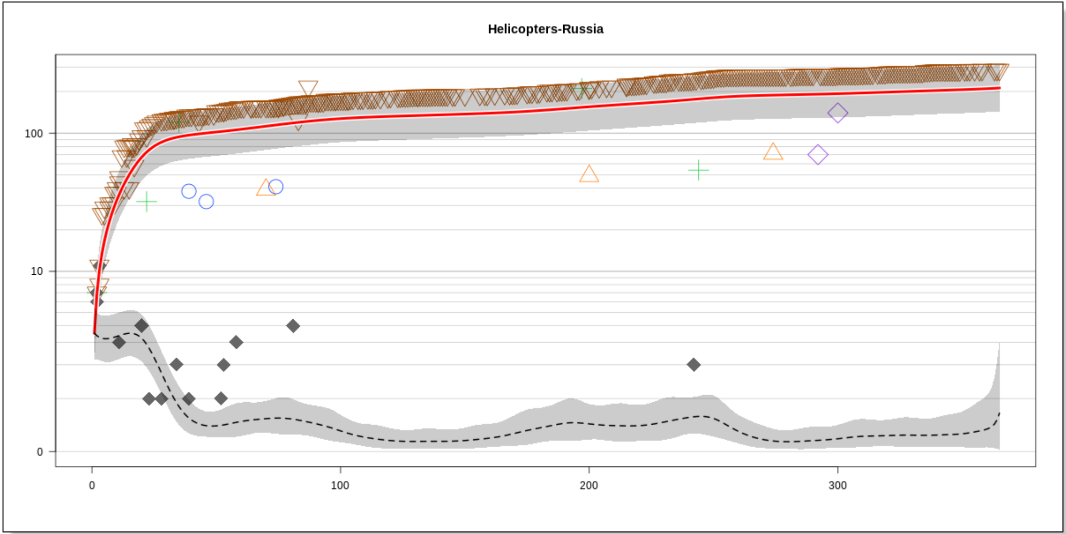

# Appendix

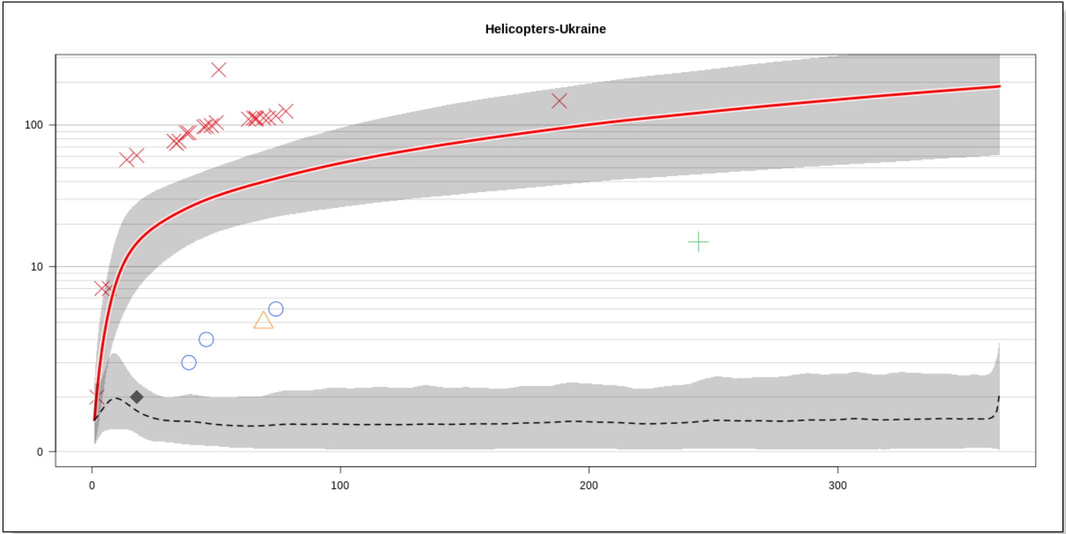

# Appendix

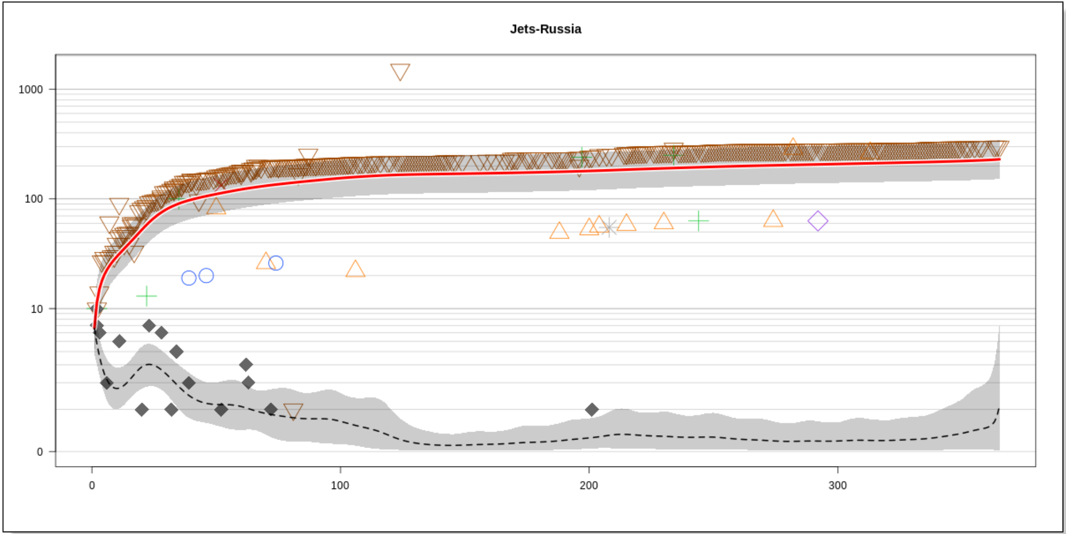

# Appendix

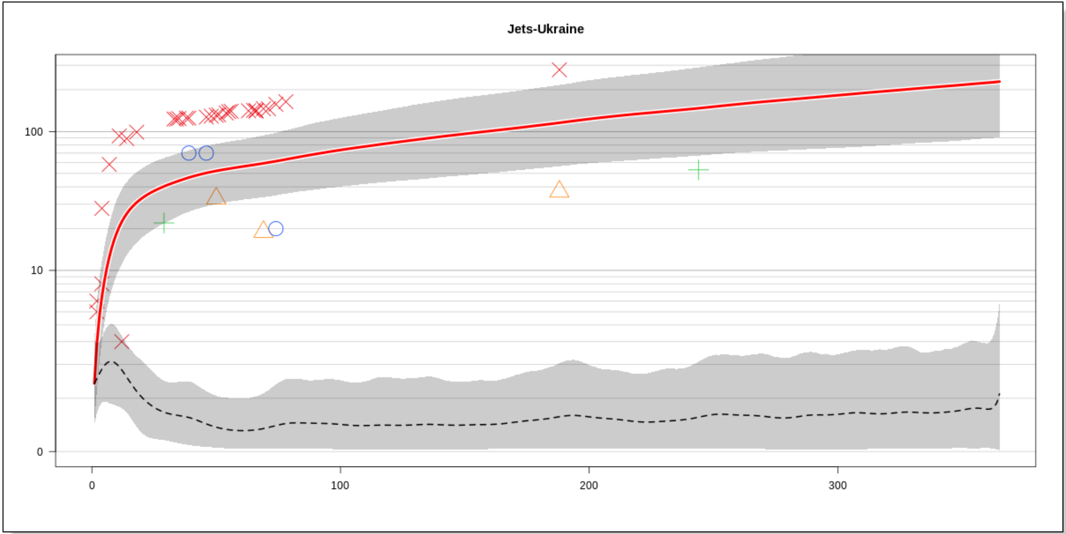

# Appendix

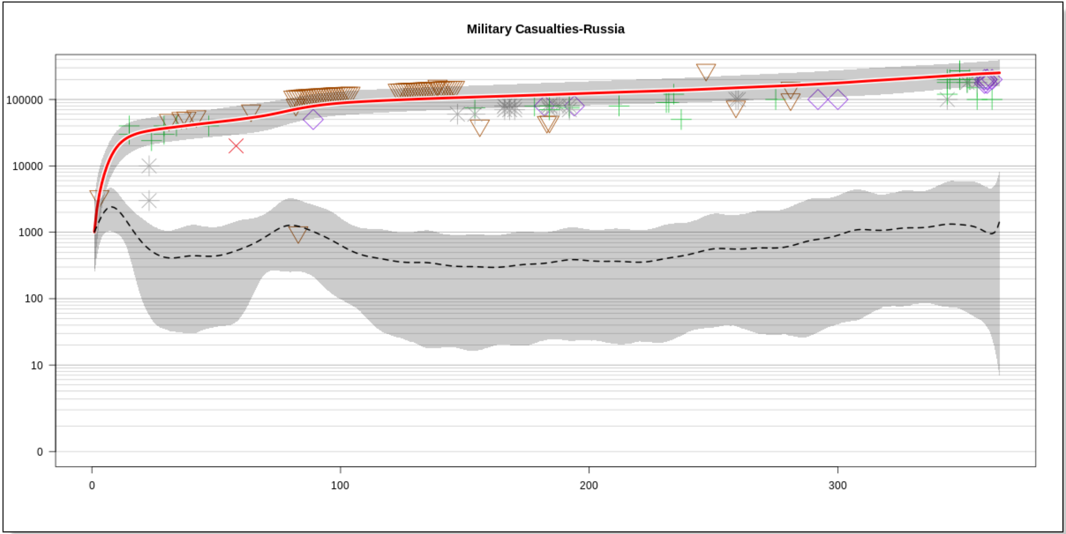

Military Casualties-Russia
100000
10000
1000
100
10
0
0
100
200
300

# Appendix

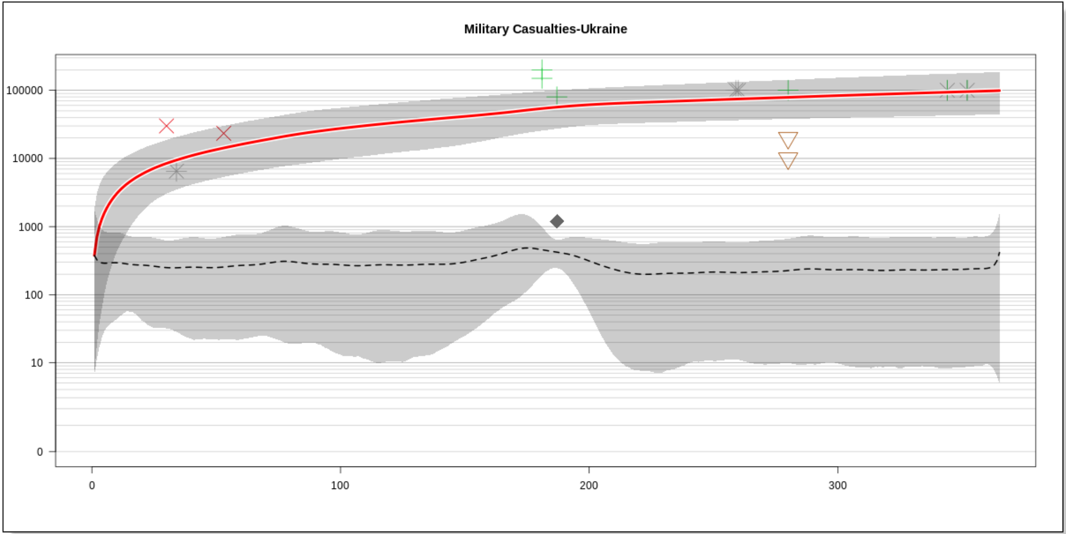

Military Casualties-Ukraine
100000
10000
1000
100
10
0
0
100
200
300

# Appendix

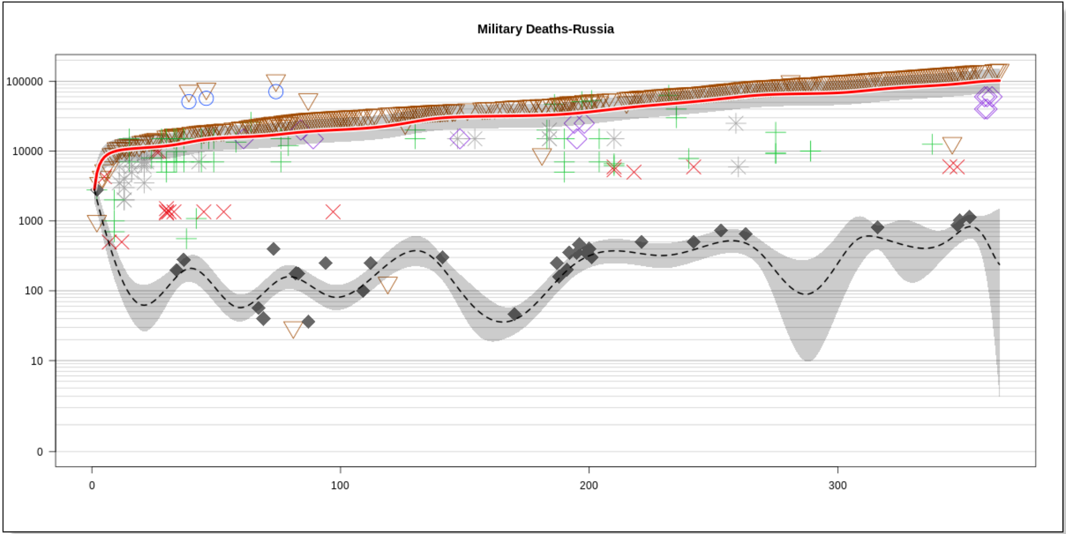

Military Deaths-Russia
100000
10000
1000
100
10
0
0
100
200
300

# Appendix

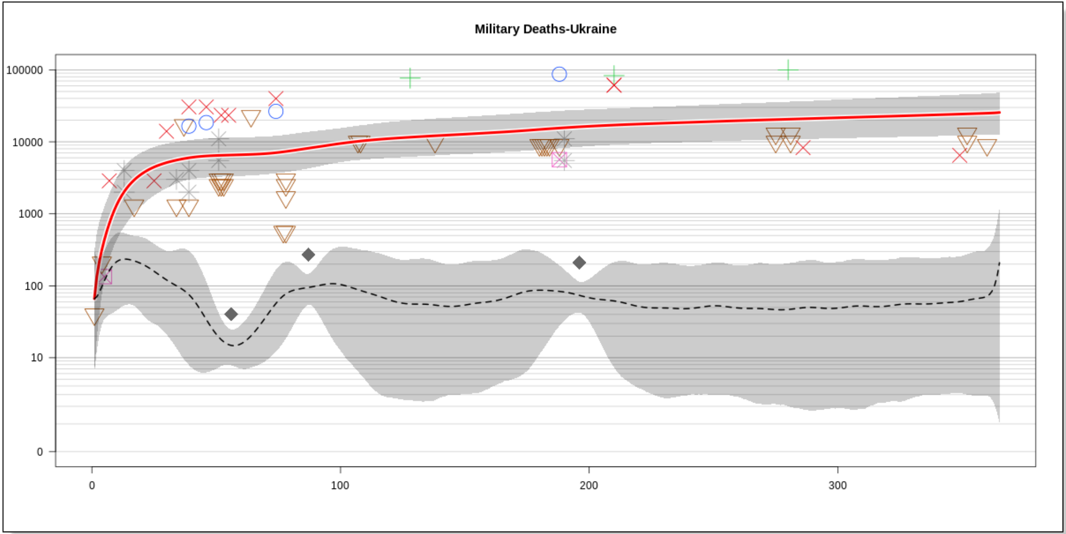

Military Deaths-Ukraine
100000
10000
1000
100
10
0
0
100
200
300

# Appendix

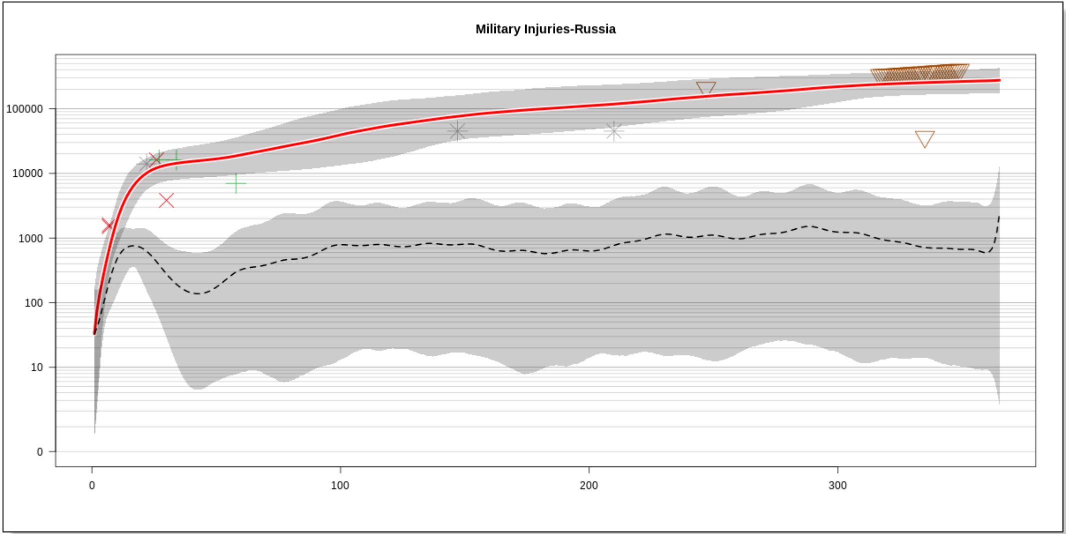
Military Injuries-Russia
100000
10000
1000
100
10
0
0
100
200
300

# Appendix

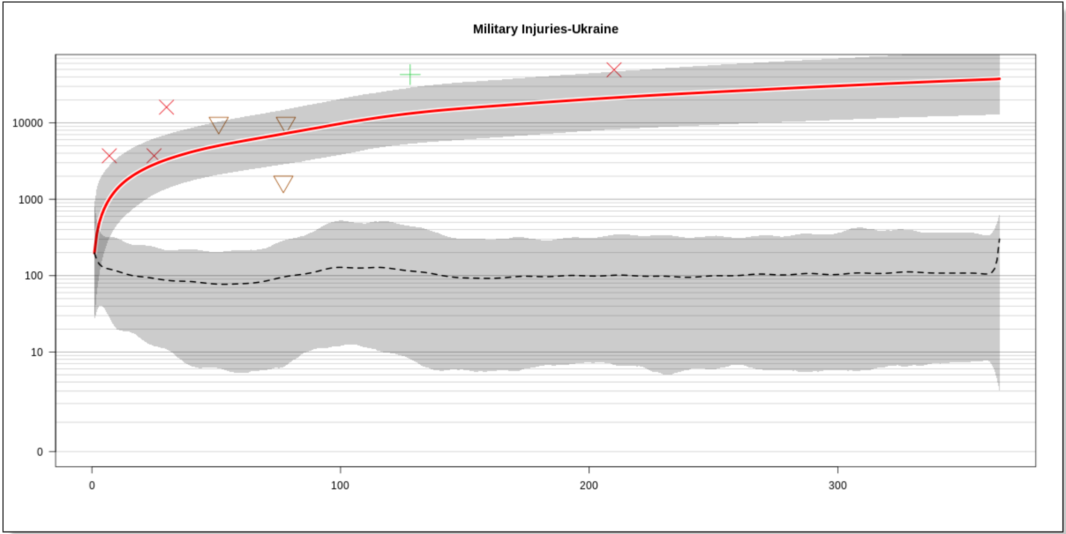

# Appendix

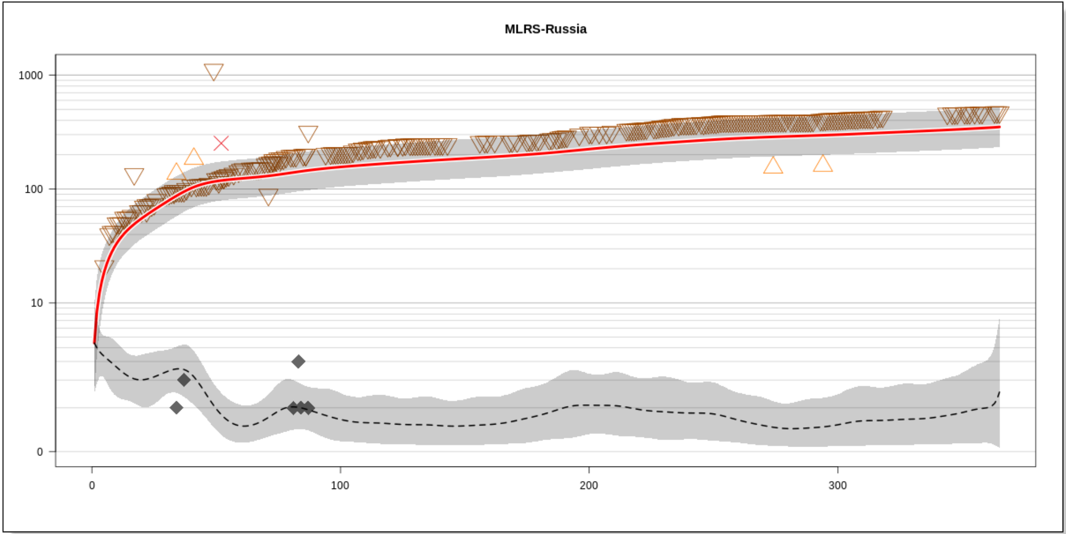

# Appendix

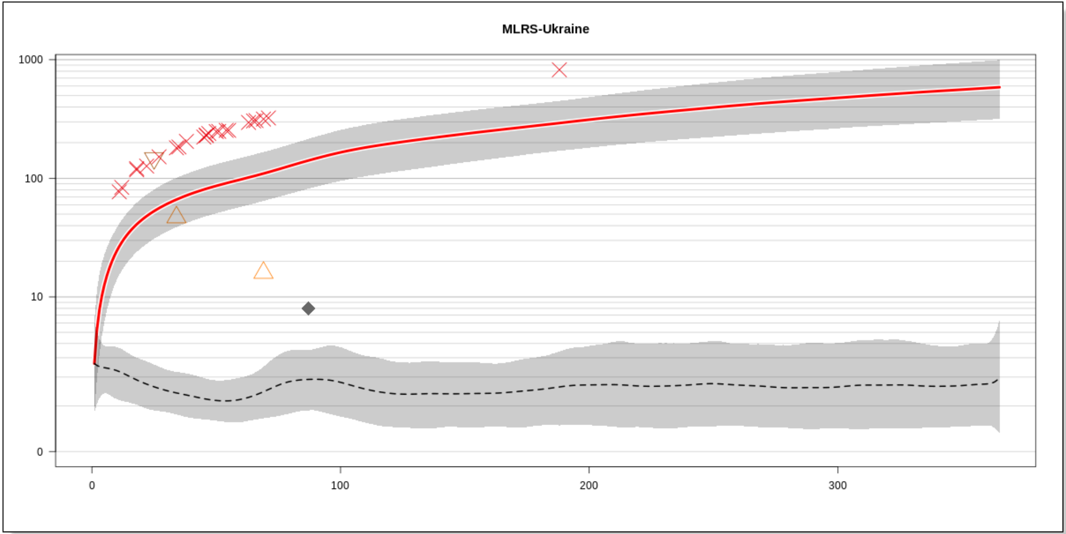

# Appendix

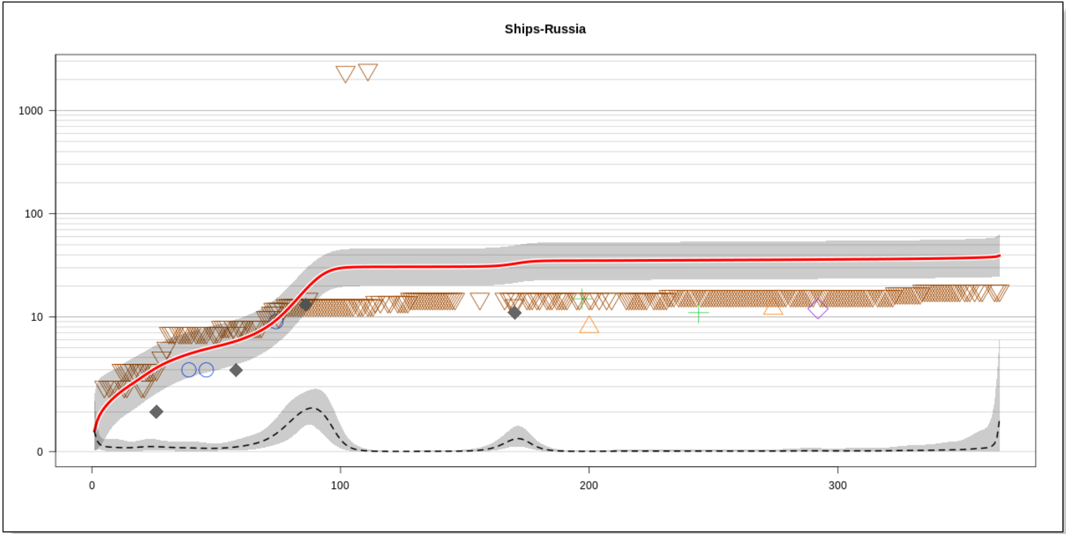

# Appendix

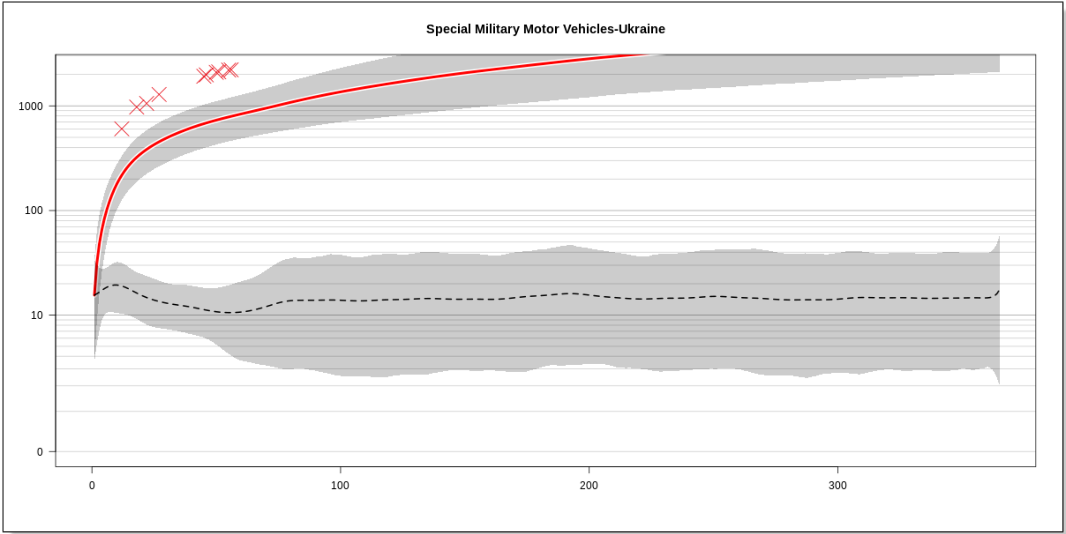

Special Military Motor Vehicles-Ukraine
1000
100
10
0
0
100
200
300

# Appendix

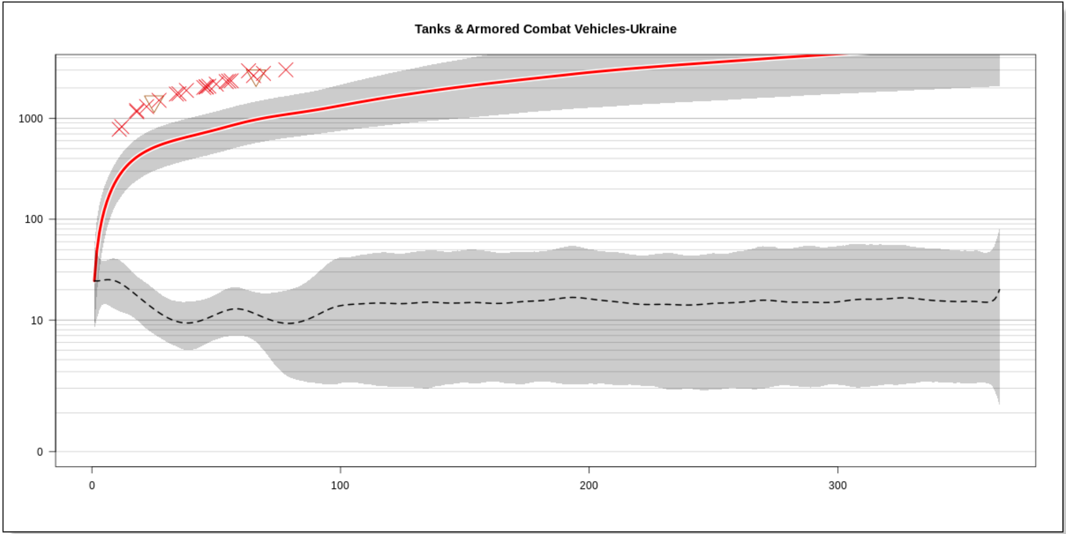
Tanks & Armored Combat Vehicles-Ukraine
1000
100
10
0
0
100
200
300

# Appendix

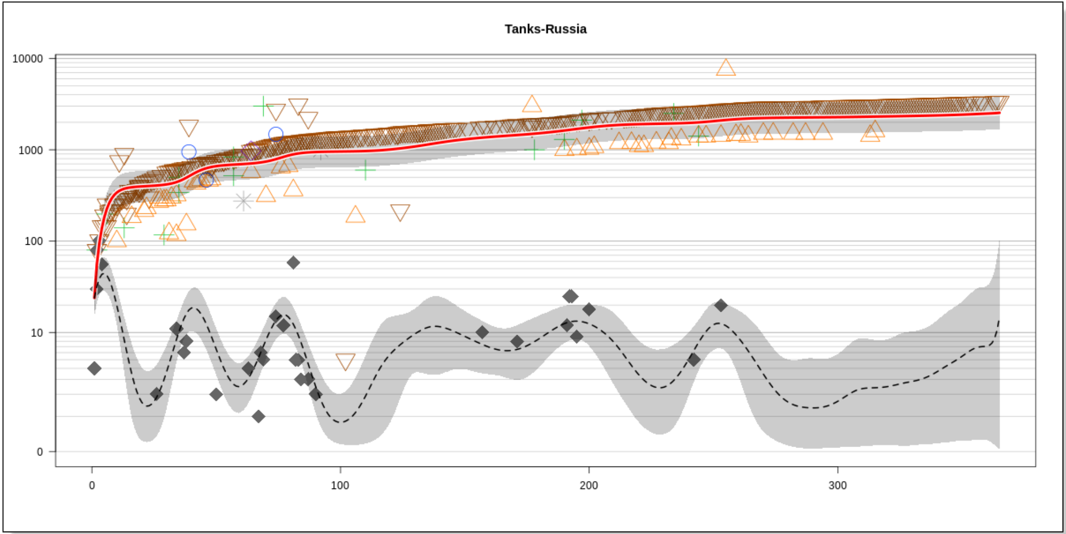

# Appendix

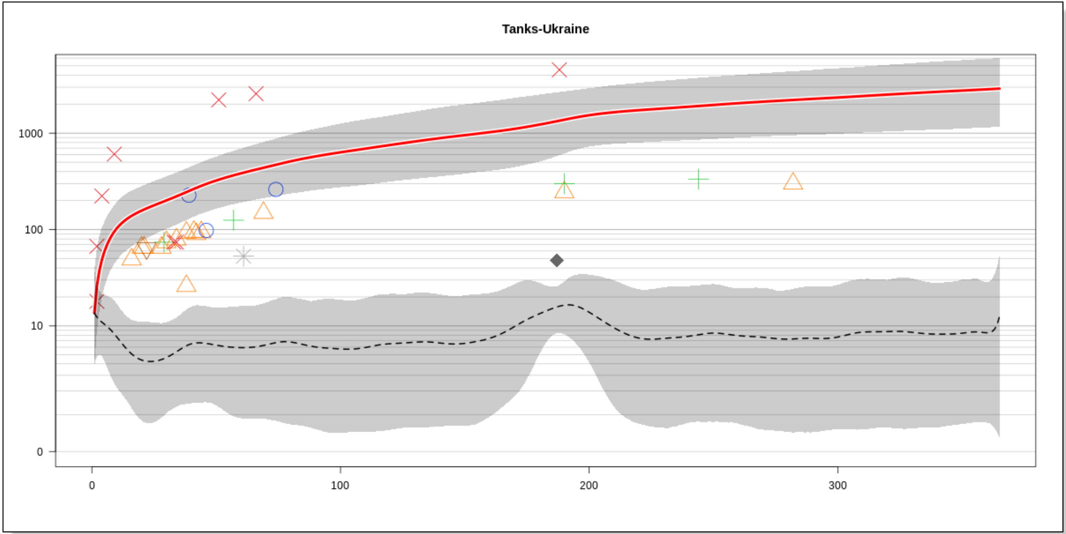

# Appendix

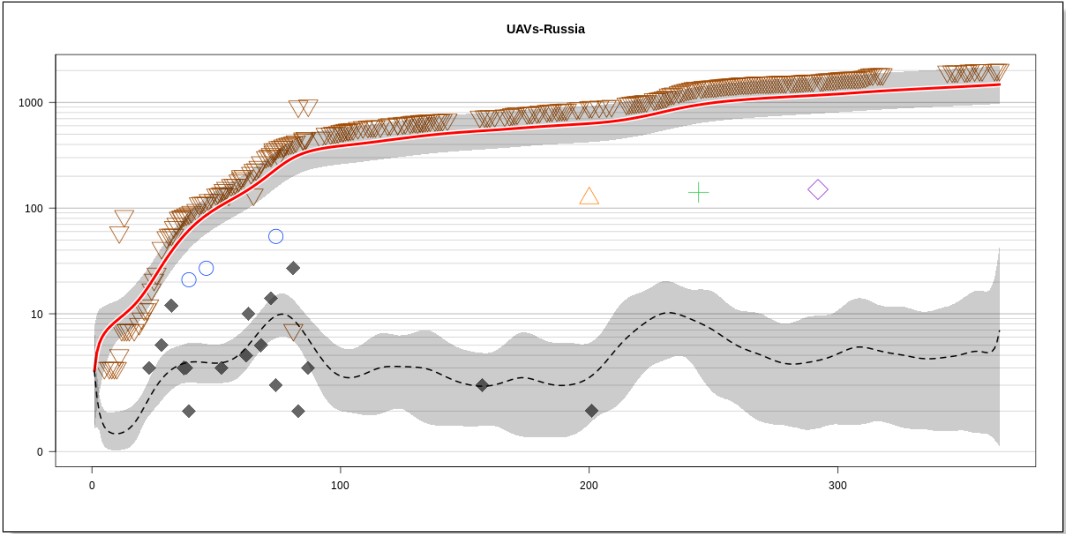

# Appendix

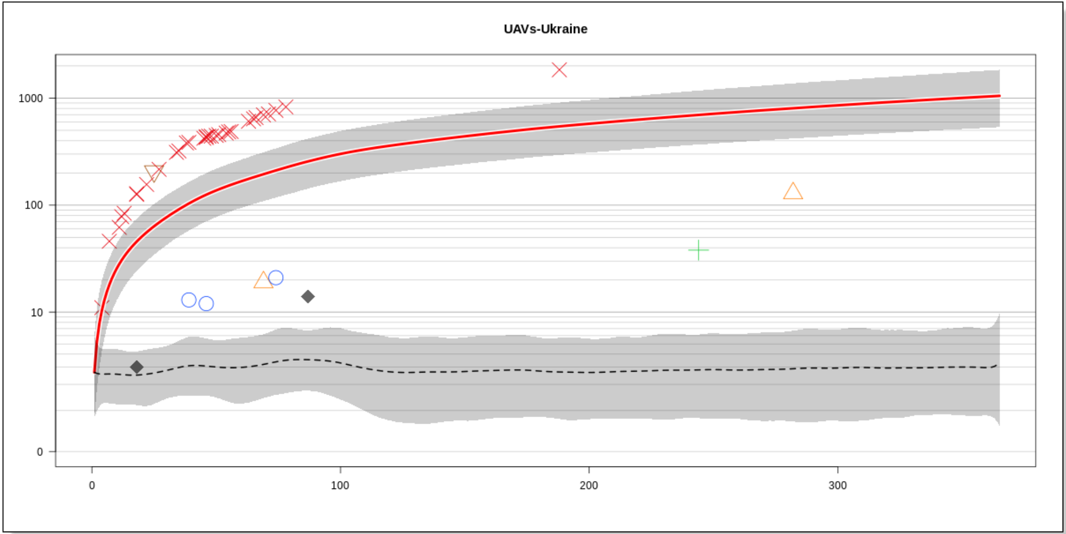

# Appendix

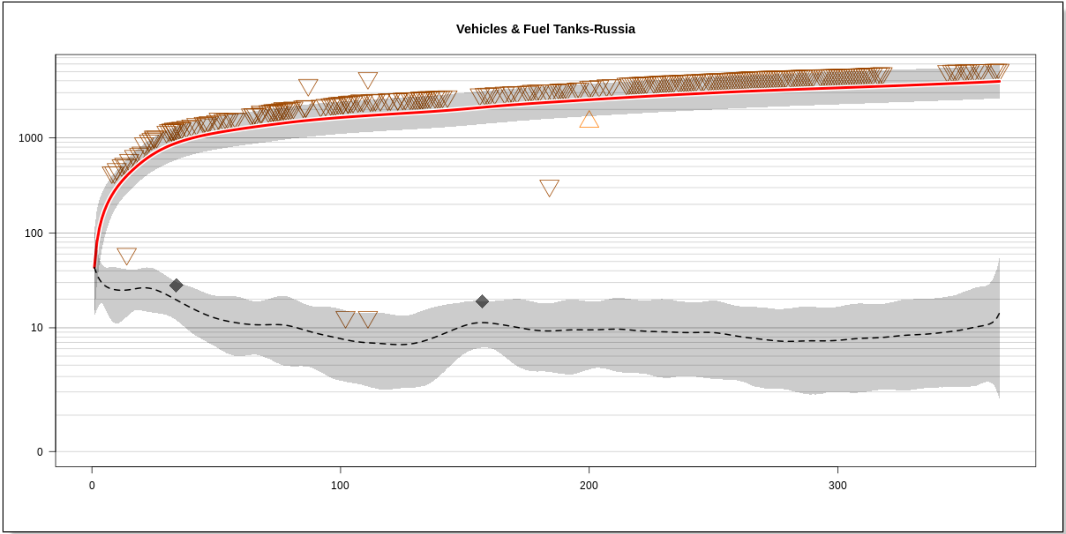

# Appendix