ETLA

Working Papers | 80

30.6.2020

Nowcasting Industrial Production Using Uncoventional Data Sources

Paolo Fornaro
Etla Economic Research
paolo.fornaro@etla.fi

Suggested citation:
Fornaro, Paolo (30.6.2020). "Nowcasting Industrial Production Using Uncoventional Data Sources".

ETLA Working Papers No 80.
http://pub.etla.fi/ETLA-Working-Papers-80.pdf

Abstract

In this work, we rely on unconventional data sources to nowcast the year-on-year growth rate of Finnish industrial production, for different industries. As predictors, we use real-time truck traffic volumes measured automatically in different geographical locations around Finland, as well as electricity consumption data. In addition to standard time-series models, we look into the adoption of machine learning techniques to compute the predictions.

We find that the use of non-typical data sources such as the volume of truck traffic is beneficial, in terms of predictive power, giving us substantial gains in nowcasting performance compared to an autoregressive model. Moreover, we find that the adoption of machine learning techniques improves substantially the accuracy of our predictions in comparison to standard linear models. While the average nowcasting errors we obtain are higher compared to the current revision errors of the official statistical institute, our nowcasts provide clear signals of the overall trend of the series and of sudden changes in growth.

Tiivistelmä

Teollisuustuotannon nykyhetken ennustaminen käyttämällä epäsovinnaisia datalähteitä

Tässä raportissa tutkitaan, miten epäsovinnaisia datalähteitä voidaan hyödyntää Suomen teollisuustuotannon vuosikasvun nykyhetken ennustamisessa (nowcasting). Käytämme sekä automaattisesti ympäri Suomen kerättyä rekkaliikenteen dataa että tietoa teollisuuden sähkönkulutuksesta, jota on saatavilla melko pian kuukauden lopun jälkeen. Perinteisten aikasarjamallien lisäksi testaamme koneoppimisen menetelmien käyttöä, kun laskemme ennusteita.

Tulokset näyttävät, että epäsovinnaiset datalähteet parantavat ennusteiden tarkkuutta verrattuna tavalliseen autoregressiviseen malliin. Lisäksi koneoppimismallien ennusteet ovat tarkempia kuin lineaaristen mallien. Mallien ennustevirheet ovat suurempia kuin Tilastokeskuksen tuottaman estimaatin tarkentumiset mutta antavat silti hyödyllistä tietoa teollisuustuotannon kasvun suunnasta.

Ph.D. (Economics) **Paolo Fornaro** is a Researcher at ETLA Economic Research.

Ph.D. (Economics) **Paolo Fornaro** on Elinkeinoelämän tutkimuslaitoksen tutkija.

Keywords: Flash Estimates, Machine Learning, Big Data, Nowcasting

Asiasanat: Pikaestimaatit, Koneoppimismallit, Massadata, Nykyhetken ennustaminen

JEL: C33, C55, E37

1 Introduction

Despite the declining manufacturing share of most Western economies, industrial production is one of the most carefully tracked macroeconomic indicators. Economists working in public and private institutions track closely how the manufacturing industry develops, both in terms of the aggregate indicator and of the sub-industries which compose it. However, as for many measures of real economic activity (such as GDP), statistical institutes publish the first estimates of industrial production with substantial delays. In Finland, the first release of industrial production is published 40 days after the end of the reference month.

In this work, we rely on unconventional data sources to provide flash estimates of the year-on-year growth rate of Finnish industrial production, for both the aggregate index and some of the sub-industries. We use real-time truck traffic volumes measured automatically in different geographical locations around Finland, as well as electricity consumption data which is available almost in real-time. In addition to standard time-series models, we look into the use of machine learning techniques to perform the predictions. We find that including information on truck traffic volumes and electricity consumption improves the nowcasting accuracy, compared to a simple autoregressive model, and that using machine learning techniques is helpful in providing better flash estimates of industrial productions, when compared to linear models. Interestingly, the inclusion of confidence indices is not beneficial in terms of nowcasting performance. However, our results are dependent on the industry considered. While we find particularly strong improvements for the aggregate industrial production index and for the paper and forest industries, our nowcasts for sub-industries such as chemicals and machinery are not more accurate than the benchmark.

Early works related to the tracking of economic conditions in real-time are Aruoba, Diebold, and Scotti (2009), for the U.S. economy, and Altissimo, Cristadoro, Forni, Lippi, and Veronese (2010) for the Euro Area, while a very recent example is found in Lewis, Mertens, and Stock (2020). In these studies, the authors develop econometric frameworks with the objective to create high-frequency indicators of economic activity. On the other hand, the nowcasting literature is interested in estimating existing economic indicators (usually quarterly GDP growth) in real-time. Few examples drawn from the nowcasting literature are Giannone, Reichlin, and Small (2008) and Evans (2005). For industrial production a very early example is provided in Ladiray and O'Brien (2003), where they use information derived from business surveys and industrial production releases from individual member states to provide faster estimates of eurozone industrial production. Brunhes-Lesage and Darné (2012) adopt both a dynamic factor model and linear regressions with automatic variable selection procedures to nowcast French industrial production, using surveys and financial data as predictors. A more recent contribution is Buono, Kapetanios, Marcellino, Mazzi, Papailias, et al. (2018), who examine the usefulness of big data in nowcasting macroeconomic series, including industrial production.

The remainder of the paper is structured as follows. We start, in Section 2, by describing our nowcasting approach, i.e. how we compute the flash estimates of industrial production and what are the models we adopt. Next, in Section 3, we go over the data used in the analysis and report the results of our exercise in Section 4. Finally, Section 5 concludes.

2 Methodological approaches

The exercise we conduct is in its essence a "horse race"; we consider a large number of modeling strategies and check their nowcasting performance for different indicators of industrial production (i.e. for different sub-industries). The models we adopt include simple time series models such as an automated ARIMA (see Hyndman and Khandakar, 2008), with and without external predictors, principal component regressions (such as in Stock and Watson, 2002), shrinkage models such as ridge regression, the least absolute shrinkage and selector operator (Lasso, see Tibshirani, 1996) and the elastic-net (see Zou and Hastie, 2005), as well as non-linear machine learning techniques. The latter include random forests, regression trees, regression splines and boosting. These machine learning techniques are implemented in `R` using the `caret` package, and an useful coverage of these models can be found in Hastie, Tibshirani, and Friedman (2009). Finally, we also consider neural networks using `kerasR`.

In addition to the techniques described above, we also compute the nowcasts using a simple combination approach, where we take the average of the flash estimates produced by individual models and give the same weight to each model. As pointed out in Stock and Watson (2004), this simple weighting strategy is effective and tends to be as competitive as more complicated ones.

Even though there are a lot of differences among the model families we consider, it is useful to give an example of the nowcasting framework we adopt. To do this, let us consider a setting where we use a linear model with lags of the dependent variable and external predictors. Our general predictive model is of the form:

$$ip_t = \beta_0 + \phi ip_{t-1} + extpred_t'\beta + \epsilon_t,$$

where ip_{t-1} is the autoregressive component and $extpred_t$ is vector of external predictors (such as confidence indices or electricity consumption). The error term in the regression is normally distributed with zero mean. Once the model is estimated, in this simple case via OLS, the prediction is given by:

$$\widehat{ip}_t = \widehat{\beta}_0 + \widehat{\phi} ip_{t-1} + extpred_t'\widehat{\beta}.$$

3 Data description

3.1 Industrial production volume

We start by describing our target variables, the monthly volume index of industrial production and some of its components. In particular, we attempt to compute a flash estimate of the working days adjusted year-on-year growth rate of these series. We plot aggregate industrial output growth in Figure 1 below. Note that the series starts in January 2007 and ends in December 2019.

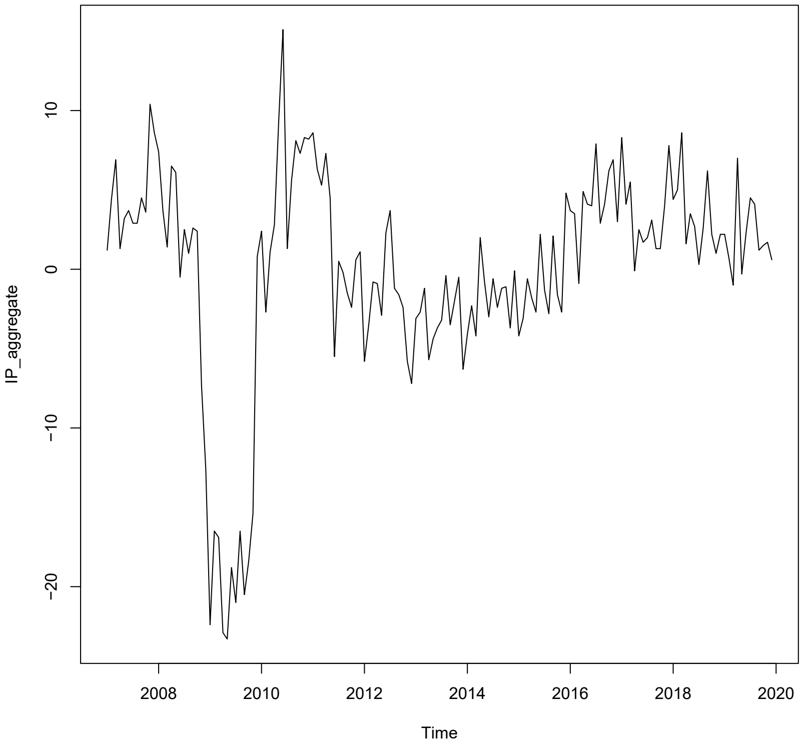

Figure 1: Year-on-year growth rate of the volume index of industrial production, working-day adjusted. The series starts in January 2007 and ends in December 2019.

As it can be seen from Figure 1, the volume index of industrial production is quite volatile, with a very large drop during the Great Recession of 2008-2009 and with a substantial rebound in the beginning of 2010. For the sake of brevity, we do not plot the rest of the target variables, but we report some basic descriptive statistics below, in Table 1.

	Aggregate	Chemicals	Forest	Machinery	Paper	Metal	Electric
Mean	0.04	2.30	-1.37	1.72	-1.27	0.93	0.34
Standard Dev.	6.74	8.43	9.45	11.66	9.48	10.39	14.25
Max.	15.1	25.4	34.3	38.3	43.4	19	45.7
Min.	-23.3	-17.4	-33.8	-34.5	-36.5	-32.9	-40.3

Table 1: Descriptive statistics for the target variables, calculated on the period running from January 2007 to December 2019

One key takeaway from Table 1 is that our target variables display considerable volatility, especially for certain sub-industries, which can results in poor nowcasting performance. This is especially true for the electric and machinery sub-industries. It is important to point out that we consider the latest release for the industrial production indices. In an optimal scenario, we would collect the data vintages for all the variables included in the exercise, but it was not possible in the current setting (especially due to the fact that we consider a number of sub-industries).

3.2 Predictors

In addition to the past values of industrial production growth, we consider a number of predictors. These include confidence indices, industrial electricity consumption and truck traffic volumes.

The confidence indicators we include are the Economic Sentiment Indicator (ESI), the Industry/Business Climate Indicator (BCI), both the aggregate index and the confidence in new orders, and the consumer confidence index. These indicators are released by the European Commission and are available around the end of the reference month, making confidence indices especially useful in terms of timeliness.

The second predictor we consider in the nowcasting model is industrial energy consumption, provided by Finnish Energy (the employers' association for the energy industry). Two features of this series are useful in our setting. Firstly, the figures provided are corrected for temperature, which is important to filter noise and pinpoint electricity consumption variation due to actual production volumes' changes. Secondly, statistics on electricity consumption are relatively timely, published usually around 20 days after the end of the reference month. This means that by using these predictors in order to produce flash estimates of industrial production we would gain 20 days in terms of publication delay, which is substantial. Of course, the actual usefulness of industrial electricity consumption for nowcasting industrial production volumes depends on the strength of the correlation between the series. We first plot, in Figure 2, the aggregate index of industrial production together with electricity consumption, then we report the figure of electricity consumption with the industrial production for the forest industry (which are strongly correlated), and finally we depict the plot of electricity consumption and the output volume for the electrical and electronics industry (which are weakly correlated).

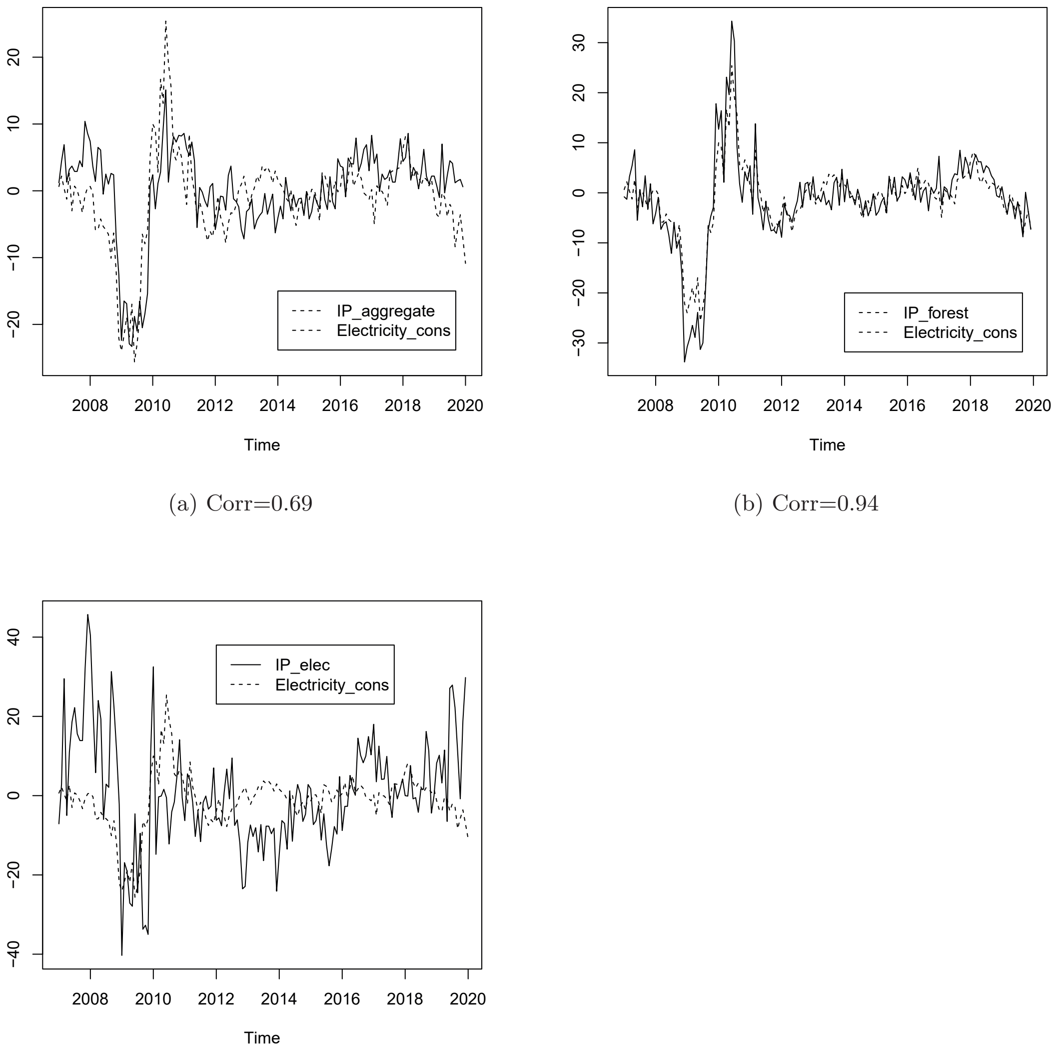

(a) Corr=0.69

(b) Corr=0.94

(c) Corr=0.18

Figure 2: Year-on-year growth rate of the volume index of industrial production, working-day adjusted, and industrial electricity consumption. Data from 2007 to December 2019.

As we can see from Figure 2, the usefulness of electricity consumption depends on the series that we want to nowcast. While industrial electricity consumption is almost perfectly correlated with the output of the forest industry, we find that the volume of production of the electric and the electronics industry is weakly correlated with electricity consumption. The aggregate index of industrial production also displays a quite strong correlation with electricity use, indicating that including this predictor in our set is probably beneficial in terms of nowcasting accuracy.

The final predictors we include in our nowcasting framework are truck traffic volumes made available by Traffic Management Finland. The data is available at the daily frequency and consists of counts of truck passages across various automatic measurement points around Finland. A more detailed discussion on how this data is treated, to obtain monthly series, is provided in Fornaro and Luomaranta (2020). In this exercise, we do not include the raw data, i.e. the traffic volumes for individual measurement points, rather we extract the first three principal components from the original data. We plot the aggregate index of industrial output and the first principal component extracted from the traffic data in Figure 3 below.

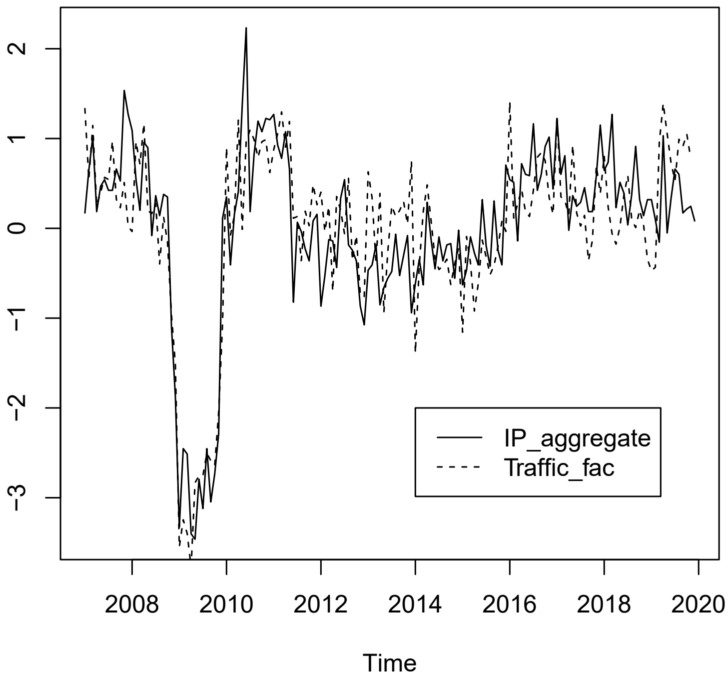

Figure 3: Year-on-year growth rate of the volume index of industrial production, working-day adjusted, and the first factor extracted from data on truck traffic volumes. The series are standardized and they cover the period from 2007 to December 2019.

From Figure 3, it is clear that truck traffic volumes are able to track well the movements in industrial production, as it is confirmed by a large correlation coefficient (0.84), and that it is a good idea to include them in the predictors' set.

4 Results

We divide the section on results by the class of models considered. For certain classes of models the number of possible specifications (e.g. due to different predictors or different parameter tuning) is large, hence to keep the discussion somewhat brief we report a subset of all models. We produce nowcasts for the period going from January 2014 until December 2019.

4.1 Time series models

The first nowcasting framework is based on standard time series techniques, such as the random walk and the automated ARIMA (our benchmark models), or the ARIMA models with external predictors. For each model we report the mean absolute error (MAE), as well as their mean error (ME).

MAE	Aggregate	Chemicals	Forest	Machinery	Paper	Metal	Electrics etc.
Random Walk	2.63	8.15	2.87	7.01	3.44	4.66	6.79
Auto. ARIMA	2.40	6.29	2.43	6.03	2.79	4.04	5.91
ARIMA+Conf.	2.47	6.81	2.61	6.37	2.97	4.68	6.31
ARIMA+Elec	2.04	6.29	2.10	5.82	2.53	4.00	5.91
ARIMA+Traf.	2.27	5.63	2.75	6.11	3.17	4.43	6.43
ARIMA+Elec.+Traf.	2.08	5.54	2.14	6.04	2.64	4.37	6.49
ARIMA+Elec.+Traf.+Conf.	2.17	5.85	2.06	6.26	2.69	4.46	6.59
ME	**Aggregate**	**Chemicals**	**Forest**	**Machinery**	**Paper**	**Metal**	**Electrics etc.**
Random Walk	0.10	0.24	−0.08	0.22	−0.04	0.04	0.27
Auto. ARIMA	0.37	1.11	0.06	1.33	-0.02	1.22	1.49
ARIMA+Conf.	0.42	0.69	0.08	0.38	0.49	0.95	0.40
ARIMA+Elec.	0.56	0.38	0.29	1.17	0.16	0.98	1.28
ARIMA+Traf.	0.33	0.54	−0.02	1.58	−0.09	1.12	1.60
ARIMA+Elec.+Traf.	0.62	0.37	0.31	1.48	0.35	1.00	1.63
ARIMA+Elec.+Traf.+Conf.	0.62	0.28	0.44	0.21	0.30	1.11	1.00

Table 2: ME and MAE for time series model. The target variables are industrial production indices and the out-of-sample period runs from January 2014 until December 2019.

The results in Table 2 give us a number of insights. Firstly, the accuracy of the flash estimates depends on the industry we consider. While for the aggregate and the forest industry the time series models provide relatively accurate predictions, the MAE for the chemicals industry, as well as the electrics and electronics industry, is very large. A second point that we can gather is that the automated ARIMA produces more precise nowcasts, compared to the random walk. However, the random walk produces the flash estimates with lowest ME (i.e. nowcasts that do not overshoot or undershoot the actual value of industrial production systematically). In terms of the usefulness of the predictors considered in the analysis, we see that confidence indices are not beneficial in terms of improving nowcasting accuracy, with the ARIMA model which includes confidence indicators providing higher nowcasting errors than the benchmark. On the other hand, truck traffic volumes and electricity consumption seem to provide useful information to the nowcasting models. However, the results in Table 2 show that models' performances depend on the target series. One thing that it is important to notice is that even the most accurate models, in terms of MAE, are not able to provide unbiased flash estimates, as it is shown by the fairly high MEs.

4.2 Machine learning techniques

The set of techniques and specifications we test in this analysis is fairly extended, so it would be impractical to report the results for each model and predictors' set. Instead we first report the MAE and ME obtained by taking a simple average of the flash estimates produced by the individual models, for each different set of predictors and industry. Note that among the predictors we include lags of both the dependent variable and of the external predictors, where we denote how many lags the considered specification includes by $Llagsy$ and $Llagsx$, where L is the number of lags for y (the dependent variable) or x (the external predictors).

Combinations							
MAE	Aggregate	Chemicals	Forest	Machinery	Paper	Metal	Electrics etc.
Elec., 5lagsy, 0lagx	1.98	6.24	2.10	5.54	2.45	3.87	5.56
Elec., 5lagsy, 3lagx	1.98	6.11	2.01	5.55	2.39	3.85	5.61
Elec., 12lagsy, 0lagx	1.96	5.81	2.02	5.30	2.22	3.73	5.71
Elec., 12lagsy, 3lagx	1.96	5.70	1.93	5.32	2.22	3.74	5.82
Elec., conf., 5lagsy, 0lagx	2.23	6.19	1.98	5.77	2.40	4.27	6.21
Elec., conf., 5lagsy, 3lagx	2.25	5.97	1.92	5.81	2.42	4.15	6.23
Elec., conf., 12lagsy, 0lagx	2.21	5.96	1.93	5.54	2.09	4.06	5.81
Elec., conf., 12lagsy, 3lagx	2.22	5.74	1.92	5.62	2.15	3.95	5.91
Elec.,traf., 5lagsy, 0lagx	1.98	6	2.19	5.68	2.50	4.17	5.70
Elec., traf., 5lagsy, 3lagx	2.03	6.03	2.20	5.97	2.50	4.33	5.70
Elec., traf., 12lagsy, 0lagx	1.96	5.72	2.17	5.47	2.31	4.06	5.73
Elec., traf., 12lagsy, 3lagx	2.03	5.85	2.26	5.84	2.40	4.42	6
Elec.,traf.,conf., 5lagsy, 0lagx	2.09	6.13	2.04	5.78	2.46	4.32	6.14
Elec., traf., conf., 5lagsy, 3lagx	2.09	5.98	2.06	5.99	2.37	4.37	6.12
Elec., traf., conf., 12lagsy, 0lagx	1.96	5.67	2.16	5.48	2.31	4.06	5.75
Elec., traf., conf., 12lagsy, 3lagx	2.01	5.87	2.23	5.86	2.39	4.39	6.01
ME	Aggregate	Chemicals	Forest	Machinery	Paper	Metal	Electrics etc.
Elec., 5lagsy, 0lagx	-0.29	-0.06	-0.67	-0.63	-0.12	-0.35	-1.23
Elec., 5lagsy, 3lagx	-0.33	-0.07	-0.61	-0.54	-0.22	-0.40	-1.06
Elec., 12lagsy, 0lagx	-0.67	-0.43	-0.86	-1.07	-0.65	-1.07	-2.58
Elec., 12lagsy, 3lagx	-0.65	-0.40	-0.80	-0.95	-0.59	-1.06	-2.63
Elec., conf., 5lagsy, 0lagx	-0.45	0.22	-0.63	-0.77	0.13	-0.51	-0.48
Elec., conf., 5lagsy, 3lagx	-0.41	0.03	-0.50	-0.96	0.10	-0.51	-0.65
Elec., conf., 12lagsy, 0lagx	-0.63	0.50	-0.84	-1.01	-0.16	-0.94	-1.96
Elec., conf., 12lagsy, 3lagx	-0.74	0.15	-0.76	-1.18	-0.05	-1.03	-2.06
Elec., traf., 5lagsy, 0lagx	-0.57	-0.19	-0.66	-0.35	-0.04	-0.76	-1.64
Elec., traf., 5lagsy, 3lagx	-0.40	0.03	-0.59	-0.29	-0.13	-0.530	-0.84
Elec., traf., 12lagsy, 0lagx	-0.93	-0.42	-1.01	-0.86	-0.66	-1.35	-2.83
Elec., traf., 12lagsy, 3lagx	-0.87	-0.34	-0.98	-0.87	-0.57	-1.17	-2.56
Elec.,traf.,conf., 5lagsy, 0lagx	-0.57	0.12	-0.57	-0.56	0.17	-0.72	-0.98
Elec., traf., conf., 5lagsy, 3lagx	-0.39	0.51	-0.53	-0.83	0.04	-0.63	-0.24
Elec., traf., conf., 12lagsy, 0lagx	-0.93	-0.45	-1.02	-0.84	-0.68	-1.34	-2.82
Elec., traf., conf., 12lagsy, 3lagx	-0.85	-0.36	-0.94	-0.88	-0.58	-1.17	-2.54

Table 3: ME and MAE for the combination of machine learning techniques and large dimensional models. The target variables are industrial production indices and the out-of-sample period runs from January 2014 until December 2019.

A few comments on Table 3. First of all, the confidence indices do not seem to be beneficial in terms of nowcasting accuracy. Moreover, while the nowcast combinations provide a slightly lower MAE, the gain compared to an ARIMA (with external predictors) model is not particularly large. One should then consider the additional computational burden associated with machine learning techniques. Finally, we see that the model combination approach produces biased nowcasts, i.e. the ME are quite high. In the next table we show the MAE for the best model for each set of predictors, where the best model is selected based on the MAE for the aggregate index of industrial volume.

MAE	Aggregate	Chemicals	Forest	Machinery	Paper	Metal	Electrics etc.
Elec.,5lagsy,0lagx(treebag)	1.89	6.08	2.44	5.93	2.74	4.09	6.37
Elec.,5lagsy,3lagx(cforest)	1.84	6.22	2.33	5.56	2.77	3.82	6.14
Elec.,12lagsy,0lagx(cforest)	1.88	5.78	2.36	5.66	2.60	3.78	6.11
Elec.,12lagsy,3lagx(cforest)	1.87	5.89	2.34	5.71	2.61	3.89	5.96
Elec.,conf.,5lagsy,0lagx (pcr)	2.19	6.59	2.10	5.48	2.54	4.29	6.67
Elec.,conf.,5lagsy,3lagx (treebag)	2.29	6.21	2.26	6.10	2.67	4.19	6.16
Elec.,conf.,12lagsy,0lagx(bstlm)	2.19	6.03	2.07	5.63	2.15	4.19	5.82
Elec.,conf.,12lagsy,3lagx(ranger)	2.21	6.19	1.86	6.12	2.24	4.16	6.06
Elec.,traf.,5lagsy,0lagx(glmnet)	1.94	5.99	2.27	5.54	2.44	4.13	5.74
Elec.,traf.,5lagsy,3lagx(ranger)	1.89	6.26	1.96	5.72	2.43	4.00	5.40
Elec.,traf.,12lagsy,0lagx(ranger)	1.95	5.88	1.83	5.68	2.24	3.99	5.99
Elec.,traf.,12lagsy,3lagx(ranger)	1.93	6.16	1.92	5.86	2.40	3.89	5.75
Elec.,traf.,conf.,5lagsy,0lagx(glmnet)	2.04	5.9	2.13	5.64	2.55	4.27	6.10
Elec.,traf.,conf.,5lagsy,3lagx(bstlm)	2.02	6.31	2.17	5.81	2.51	4.29	6.29
Elec.,traf.,conf.,12lagsy,0lagx(cforest)	1.98	5.71	2.32	5.81	2.540	4.04	6.00
Elec.,traf.,conf.,12lagsy,3lagx(ranger)	1.91	6.08	1.88	5.88	2.38	3.83	5.75

Table 4: MAE for the machine learning technique yielding the most accurate nowcasts for a given set of predictors. The target variables are industrial production indices and the out-of-sample period runs from January 2014 until December 2019.

Table 4 confirms that our ability to produce accurate (relatively) flash estimates depends on the industry considered. While we do quite a good job in producing nowcasts for the aggregate index of industrial production, industries like chemicals, machinery and electrics show large nowcasting errors. However, we can find some general insights: as before, confidence indices do not seem to be beneficial in our context, while industrial electricity consumption is the most useful predictors to predict the aggregate index of industrial production. Traffic volumes do not seem, surprisingly, to be helpful in predicting total industrial production but they provide quite accurate nowcasts for the forest industry. Finally, our results indicate that random forests techniques do pretty well in our context, beating also the combination approach.

4.3 Neural networks

The final models we consider in this study are neural networks. While these techniques have been show to be effective in numerous predictive contexts, they suffer from the abundance of hyperparameters that need to be set by the user. For example, using the `kerasR` interface, we need to decide how many layers the neural network has, how many nodes we should include in each layer, what kind of optimization procedure we use and how much weight decay (how much shrinkage) we put in the model, just to name a few. Of course, on top of this we also need to consider the possible sets of predictors, in terms of lags of the dependent variable and of the external predictors, as well as which of the latter we include.

While we hope that the predictive performance of the neural networks is relatively robust to different hyperparameters specifications, this is not always true. In an optimal scenario, we would use a cross-validation approach where we consider a very large number of specifications and combinations of hyperparameters, to then select the ones which offer the best predictive performance on a validation set. However, this approach requires a lot of computational power, usually offered by cloud computing services. Given the exploratory purpose of this work, we cannot do this, so we limit ourselves to consider a relatively small number of specifications for each set of predictors. We report only the MAE and ME for the specification giving the best MAE for the aggregate index of industrial production, and for a given set of predictors, in Table 5 below.[1]

Predictors	MAE	ME
Elec.,5lagsy,0lagx	1.91	0.05
Elec.,5lagsy,3lagx	2.05	-0.01
Elec.,12lagsy,0lagx	2	-0.47
Elec.,12lagsy,3lagx	2.04	-0.27
Elec.,conf.,traf.,5lagsy,0lagx	1.91	-0.08
Elec.,conf.,traf.,5lagsy,3lagx	2.05	-0.48
Elec.,conf., traf., 12lagsy,0lagx	1.99	-0.41
Elec.,conf., traf, 12lagsy,3lagx	2.14	-0.69
Elec.,traf.,5lagsy,0lagx	1.88	-0.1
Elec.,traf.,5lagsy,3lagx	1.91	$-0,19$
Elec.,traf.,12lagsy,0lagx	1.94	-0.38
Elec.,traf.,12lagsy,3lagx	1.92	-0.41

Table 5: ME and MAE for the nowcasts of neural network models. For each set of predictors considered, we report the results of the neural network specification yielding the lowest MAE. The target variable is the aggregate index industrial production and the out-of-sample period runs from January 2014 until December 2019.

We can draw a few insights, based on the results of Table 5. It seems that considering

[1]Results for all combinations of hyperparameters considered are available upon request

longer lags of the target variable in the set of predictors leads to a worse nowcasting performance, while including 5 lags of industrial production in the predictors improves the nowcasts (we have tried specifications with no or very few lags, and these led to worse results). The second point worth highlighting is that including lags of the external predictors (e.g. of electricity consumption) does not help the nowcasts. Finally, including both electricity consumption and traffic volumes in the set of predictors improves the nowcasting performance of our models, while confidence indices are not helpful.

The overall predictive ability of the neural network models is fairly good, especially in the light of the fact that the best estimates obtained from these techniques are unbiased (i.e. they have a low mean error). In particular, the best model reported in Table 5 has relatively low MAE (as low as the one produced by the best machine learning technique) and low ME. On the flipside, we find that the neural network results are highly dependent on the hyperparameter we chose, hence the selection of the best model to adopt going forward can be time consuming.

4.4 Practical usefulness of the nowcasts

The results we have obtained so far indicate that the use of data sources like electricity consumption and truck traffic volumes are useful in nowcasting industrial production. In particular, adding external predictors in our models improve the nowcasting performance, compared to an automated ARIMA approach. Moreover, the use of non-linear techniques, such as random forests and neural network yields further improvements. However, especially for certain industries, the prediction error remains large.

Looking at past revisions of the volume index of industrial production (total) by Statistics Finland, we find a mean absolute error of around 0.8 percentage point, while our best models give a MAE ranging from 1.84 (conditional random forest) to 1.88 (neural network). This result indicate that our nowcasts are of relative usefulness, when considering individual months. We do gain 20 to 25 days in publication lag, but the increase in revision error is substantial.

However, early estimates can be useful in terms of signaling the general trend of the target indicator, even though they are not able to predict accurately individual months. To verify whether our nowcasts offer a meaningful signal of the evolution of aggregate industrial output, we look at the 3-months moving average for both the nowcasts and the target indicator. For this exercise, we consider a time period going from March 2014 to March 2020. The nowcasts we use here are the ones produced by the neural network giving the best MAE in Table 5, i.e. the model with 5 lags of industrial output, electricity consumption and traffic volumes. Moreover, we also plot the nowcasts from a neural network model where we include only one lag of the target variable in the predictors' set, to see whether a smaller degree of persistence in the model helps in predicting sudden changes in growth trends.

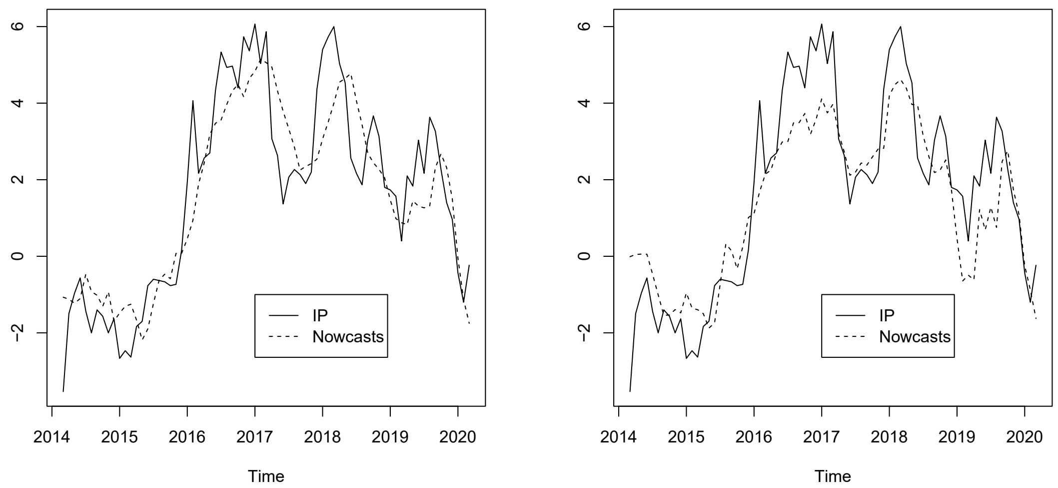

(a) MAE=0.93, ME=0.11, Corr=0.89

(b) MAE=0.99, ME=0.30, Corr=0.89

Figure 4: Year-on-year growth rate of the volume index of industrial production, working-day adjusted, and nowcasts produced by neural network models. The model including 5 lags of the target variables are reported in subfigure (a), while the nowcasts in subfigure (b) include the first lag of industrial production among the predictors. Series are smoothed using a 3-months moving average and the period considered runs from January 2014 until March 2020.

The plot (a) in Figure 4 indicates that our best nowcasts, in terms of MAE, are strongly correlated with the target series, once smoothed, however they seem to show a marked lag compared to actual industrial production. For example, the predictions of our model struggle to signal the steep increase in growth at the beginning of 2016, or the one appearing during the middle of 2017. The same goes for the large drop around the mid of 2018. However, both models seem to capture the sudden (albeit relatively) moderate drop in industrial production in February, at the start of the Covid-19 crisis.

We report the nowcasts obtained from the model using only one lag of industrial production in the predictors' set in subfigure (b). While we find that these nowcasts are slightly worse compared to the ones depicted in (a), in terms of MAE and ME, they do display a very strong correlation with the target series. Moreover, and more importantly, the model with fewer lags seems to be able to track sudden changes in growth better than the most accurate model. This can be seen in the upswing in growth during the beginning of 2016 or in the drop at the beginning of 2017, where the more parsimonious model produces nowcasts which do no lag as much as the ones described above. This insight can be the focus of future research, where nowcasts are targeted at signaling sudden changes in the series behavior and on directional predictions.

5 Conclusions

The main objective of this exercise was to verify the usefulness somewhat atypical data sources in computing flash estimates of industrial production. In particular, our main focus was on determining how beneficial can electricity consumption and truck traffic volumes be in nowcasting manufacturing activity.

We find that industrial electricity consumption and, to a lesser extent, truck traffic volumes are useful in predicting industrial production in real time. However, the inclusion of confidence indices in the models does not bring much value. Furthermore, we find that typical time series approaches, such as ARIMA models, are outperformed by non-linear techniques such as neural networks or random forests. These considerations hold for certain target indicators, such as aggregate industrial production and the index of production volumes of the forest industry.

It is important to highlight that the nowcast errors exceed the current revision errors of Statistics Finland by quite a large margin (roughly a difference in the mean absolute error of 1 percentage point, in favor of the latter), weighting down the actual usefulness of these nowcasts in a real-time setting, at least when considering individual months' estimates. On the other hand, the nowcasts are quite useful in signaling abrupt changes in industrial production growth, where a reduction in publication delays can be quite important.

References

Filippo Altissimo, Riccardo Cristadoro, Mario Forni, Marco Lippi, and Giovanni Veronese. New Eurocoin: Tracking Economic Growth in Real Time. *The Review of Economics and Statistics*, 92(4):1024–1034, November 2010.

S. Boragan Aruoba, Francis X. Diebold, and Chiara Scotti. Real-Time Measurement of Business Conditions. *Journal of Business & Economic Statistics*, 27(4):417–427, 2009.

Véronique Brunhes-Lesage and Olivier Darné. Nowcasting the french index of industrial production: A comparison from bridge and factor models. *Economic Modelling*, 29 (6):2174–2182, 2012.

Dario Buono, George Kapetanios, Massimiliano Marcellino, Gianluigi Mazzi, Fotis Papailias, et al. Big data econometrics: Now casting and early estimates. *Milan, Bocconi University, Baffi-Carefin centre working paper*, 82, 2018.

Martin D. D. Evans. Where Are We Now? Real-Time Estimates of the Macroeconomy. *International Journal of Central Banking*, 1(2), September 2005.

Paolo Fornaro and Henri Luomaranta. Nowcasting finnish real economic activity: a machine learning approach. *Empirical Economics*, 58(1):55–71, 2020.

Domenico Giannone, Lucrezia Reichlin, and David Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55 (4):665–676, May 2008.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 edition, 2009.

Rob Hyndman and Yeasmin Khandakar. Automatic time series forecasting: The forecast package for r. *Journal of Statistical Software, Articles*, 27(3):1–22, 2008.

Dominique Ladiray and Derry O'Brien. Nowcasting Eurozone Industrial Production . Technical report, 2003.

Daniel J. Lewis, Karel Mertens, and James H. Stock. Monitoring Real Activity in Real Time: The Weekly Economic Index. Liberty Street Economics 20200330b, Federal Reserve Bank of New York, March 2020.

James H. Stock and Mark W. Watson. Forecasting Using Principal Components From a Large Number of Predictors. *Journal of the American Statistical Association*, 97: 1167–1179, December 2002.

James H. Stock and Mark W. Watson. Combination forecasts of output growth in a seven-country data set. *Journal of Forecasting*, 23(6):405–430, 2004.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 67(2): 301–320, 2005. ISSN 1369-7412.

Elinkeinoelämän tutkimuslaitos

ETLA Economic Research

ISSN-L 2323-2420
ISSN 2323-2420 (print)
ISSN 2323-2439 (pdf)

Publisher: Taloustieto Oy

Tel. +358-9-609 900
www.etla.fi
firstname.lastname@etla.fi

Arkadiankatu 23 B
FIN-00100 Helsinki