Predictive Power of Dynamic Risk Factors in the Finnish Risk and Needs Assessment Form
Compared to Static Predictors

Benny Salo¹, Toni Laaksonen¹, Pekka Santtila^{1,2}

¹Åbo Akademi University, Faculty of Arts, Psychology and Theology

²New York University Shanghai, Faculty of Arts and Sciences

Author Note

Corresponding author: Benny Salo, benny.salo@abo.fi. Faculty of Arts, Psychology and Theology, Åbo Akademi, Domkyrkotorget 3, 20500 Åbo.

This work was supported by the Criminal Sanctions Agency in Finland. It was also supported by personal funding: for the first author by the National Doctoral Programme of Psychology in Finland, the Finnish Cultural Foundation, the Åbo Akademi University Foundation, Svensk-Österbottniska Samfundet and Waldemar von Frenckells stiftelse; and for the third author by the Academy of Finland Project 287800.

Analyses are documented and results are provided in R package form on <https://github.com/bennysalo/predict-recidivism>.

This is a revised version of a preprint originally published on February 19, 2018. This version is under peer-review as of May 4, 2018.

Abstract

We estimated the predictive power of the dynamic items in the Finnish Risk and Needs Assessment Form (RITA), assessed by case-workers, for predicting recidivism. These 52 items were compared to static predictors including crime committed, prison history, and age. We used two machine learning methods (elastic net and random forest) for this purpose and compared them with logistic regression. Participants were 746 men that had, and 746 that had not, reoffended during matched follow-up periods from 0.5 to 5.8 years. Both RITA-items and static predictors predicted general and violent recidivism well ($AUC = .73 - .79$), but combining them increased discrimination only slightly ($\Delta AUC = 0.01 - 0.02$) over static predictors alone. Calibration was good for all models. We argue that the results show strong potential for the RITA-items but that development is best focused on improving usability for identifying treatment targets and for updating risk assessments.

Keywords: risk and needs assessment, dynamic risk factors, recidivism, machine learning, RITA, Finland

Predictive Power of Dynamic Risk Factors in the Finnish Risk and Needs Assessment Form
Compared to Static Predictors

Risk assessment in correctional settings has two important, and partly competing, roles: to assess the risk of new crimes and thereby assign prisoners to the appropriate level of supervision, and to identify needs of the prisoner to provide treatments that can facilitate successful reintegration into society (Monahan & Skeem, 2016). It is important to validate risk assessment instruments in the setting they are to be used (Gottfredson & Moriarty, 2006). For an instrument to be informative for decisions regarding supervision level it has to show predictive validity for the unwanted events that assigning higher security levels aims to avoid. To the extent needs are identified as treatment targets, they not only have to predict future success of reintegration but also explain the level of future success (i.e. construct validity). A textbook case of thorough instrument development would start with creating theory-driven measurements, then establishing the construct validity for these measurements, and finally estimating their predictive validity. However, there are often data, with measurements of differing levels of theoretical justification and construct validity, readily available for analysis. In these cases, as exemplified in the present study, we argue that an explicit examination of the predictive validity is a good place to start. We aimed to show that two machine learning methods (*elastic net logistic regression* and *random forest*) are well suited for this task and to illustrate the insights such an examination can provide about the role of the particular set of so called *dynamic risk factors* presented in the present study.

In risk assessment, a distinction is made between *static risk factors*, that cannot be changed through interventions (e.g. number of previous sentences or age), and dynamic risk

factors that can change and be changed during the prison sentence (e.g. alcohol use or employment problems). Instruments that assess dynamic risk factors have potential benefits over only assessing static predictors. They hold the promise of being able to identify treatment targets at the same time as the current risk level is assessed. Variables that can change can also be used to update the risk assessment. The *Finnish Risk and Needs Assessment Form (Riski- ja tarvearvio, RITA)* is an example of an assessment tool that puts focus on risk factors that are at least potentially dynamic. In the present study we compared the predictive power of the dynamic items in RITA to available static predictors.

The Value of Quantifying Predictive Power

Starting from predictive validity, instead of construct validity, risks resulting in multi-dimensional composite constructs rather than construct valid and causal constructs (Cording, Beggs Christofferson, & Grace, 2016; Ward, 2016). However, focusing on predictive power has several scientific uses, among them the aim to quantify the extent to which an outcome is predictable with the information at hand (Shmueli, 2010). This can in the long run help develop construct valid causal constructs. We should seek to better understand variables with high predictive validity, on the other hand, variables that do not predict a relevant outcome should be abandoned or redesigned.

It is worth noting that *predicting* and *explaining* a phenomenon, such as recidivism, are tasks that put different demands on both the variables and the statistical methods used (Shmueli, 2010). Regarding statistical methods, the goals of accurate prediction and of interpretable explanation may be incompatible in the sense that maximizing one often entails compromises on the other (Kuhn & Johnson, 2013). When we try to explain how recidivism can be predicted we focus on the function that links the predictors to the outcome. This function is likely to be

complex and we often need to simplify it into a manageable statistical model to be able to describe it. An alternative is to ease our requirement to understand our model in favor of maximizing how well it duplicates the results of the true function (Breiman, 2001b). This allows us to more explicitly test the predictive power of dynamic risk factors compared to static predictors before introducing constraints that will improve interpretability but reduce predictive power. It can be argued that treating the relationship between predictors and outcome as a *black box* is unsatisfactory (Zeng, Ustun, & Rudin, 2017), but we can ultimately choose to use a more transparent or more easily applied scoring technique. In that case, we are interested in how much information is lost (if any) in terms of predictive power in this simplified application.

Risk Assessment in Finnish Prisons

In Finland, in accordance with the Finnish Imprisonment Act (*Vankeuslaki*, 767, 2005), a plan is drawn up in the beginning of the sentence for how the term should be served. The plan includes placement and planned activities during the sentence. In this sentence plan, available information about previous sentences, working and functional ability, criminality, and other circumstances are considered. The plan is drawn up in one of the regional assessment and allocation units and amended in the prison unit where the person is eventually placed. In some cases (typically for persons with sentences longer than one year) the assessment unit uses a structured assessment form, here termed the *Finnish Risk and Needs Assessment Form* (*Riski- ja tarvearvio*, RITA). It is based on the *Offender Assessment System* (OASys) used in England and Wales (Debidin, 2009). It was translated and adapted for the Finnish circumstances in 2004 (Lilja, 2014) and taken into common use by the Criminal Sanctions Agency in Finland in 2006.

The RITA contains two opening sections on the details of the committed crimes for the current sentence and eight sections covering dynamic variables. These sections are:

Accommodation and managing activities of daily living; Income and managing financial situation; Education, employment and skills supporting education and employment; Social bonds and life-style, Alcohol use; Drug use; Thinking and behavior; and Attitudes.

The form has no formal status in any process of correctional decision-making. For example, there are no regulations as to what kind of RITA profile an inmate should have in order to be placed in an open institution (Criminal Sanctions Agency in Finland, 2004). It can be argued that the individual benefits from needs being addressed regardless of whether or not it reduces rates of recidivism, but assessing dynamic risk factors holds additional potential value.

Value of Dynamic Risk Factors

A particularly important reason for considering dynamic risk factors is that since they can change, they can be used to update assessments of risk (Clarke, Peterson-Badali, & Skilling, 2017; De Vries Robbé, de Vogel, Douglas, & Nijman, 2015; Howard & Dixon, 2013; Lewis, Olver, & Wong, 2013; Olver, Beggs Christofferson, Grace, & Wong, 2014). If change occurs due to intervention, and if this change can be linked to changes in recidivism rates, then assessing dynamic risk factors can even help us establish causal explanations for recidivism (Cording et al., 2016; Ward & Beech, 2015). Repeated measurement of the same variables is needed for this potential to come into play, and for now this is not a feature of the Finnish risk assessment system. The RITA could be characterized as a “third-generation risk assessment instrument” (Andrews, Bonta, & Wormith, 2006; Campbell, French, & Gendreau, 2009) that includes dynamic variables but is not integrated into assessment of change, or success in interventions. Without evidence that the variables actually change or that this change has any effect on probability of recidivism they may cautiously be called *variable markers* (Kraemer, 2003) or *potentially dynamic risk factors* (Hanson & Harris, 2000).

However, risk assessments are needed, for early decisions, before change has had the chance to occur. Basing an initial assessment on dynamic factors can simultaneously facilitate decisions regarding placement and regarding planned interventions. To justify a higher cost of assessing dynamic factors they should be truly informative, and preferably show incremental predictive power over easily assessed static predictors.

Predictive Validity of Risk Assessment Instruments

There is ample evidence for actuarial assessment that use standardized questions and scoring schemes generally having higher predictive power than clinical judgement that use neither (Ægisdóttir et al., 2006; Andrews et al., 1990; Dawes, Faust, & Meehl, 1989; Grove, Zald, Lebow, Snitz, & Nelson, 2000; Quinsey, 2009). However, among these structured assessments, previous research does not show a clear advantage of a particular set of predictors. In meta-analyses, risk assessment instruments that include, and those that do not include, dynamic risk factors showed similar predictive performance (Campbell et al., 2009; Yang, Wong, & Coid, 2010). Furthermore there seem to be considerable overlap in the predictive information that items within an instrument provide. Kroner, Mills, and Reddon (2005) randomly picked 13 items, to construct new measures, from four well-established risk assessment instruments: *Psychopathy Checklist – Revised* (PCL–R; Hare, 2003), LSI–R (Andrews, Bonta, & Wormith, 2004), VRAG (Quinsey et al., 2006) and the *General Statistical Information on Recidivism* (GSIR; Nuffield, 1982). The randomly assembled instruments predicted recidivism at the same level as the original instruments. Coid et al. (2011) showed that only a subset of items within PCL-R, VRAG and *Historical-Clinical-Risk Management-20* (HCR-20, Webster, Douglas, Eaves, & Hart, 1997) is needed to account for the predictive validity of the instruments.

There is clearer evidence that the methodology of a validation study plays a role in the predictive power that a study finds. Among others, these factors include reliability of predictor measures, quality of the outcome measure, base-rate of recidivism, and length of follow-up (Andrews et al., 2011; Olver, Stockdale, & Wormith, 2014; Yang et al., 2010). More careful research methodology in studies where the instruments are developed may explain an authorship effect found by Singh, Grann, and Fazel (2013) where instruments perform better in studies where instrument developers are involved. This raises the question of whether risk assessment instruments, those with dynamic predictors in particular, perform as in applied as in research settings (Cording et al., 2016). Flores, Lowenkamp, Holsinger, and Latessa, (2006) showed that LSI-R predicted recidivism a lot better when the assessment was done by a person trained specifically in the use of LSI-R and when it was done by a person with more experience.

A research result, that does not show this type of effect, hints at parallel explanation. Jones, Brown, and Zamble (2010) found that parole-officer and researcher ratings resulted in models with similar predictive validity. However, the models in their study, derived from stepwise Cox regression, were allowed to be different for different types of raters. Thus, it is possible that any differences in rating quality was offset by customized weighting of the model predictors. Weighting that is customized to the setting where the instrument is used may have positive effects on the predictive validity even if changed scoring rules might hurt construct validity.

Methodological Considerations for Examining Predictive Power

There has been considerable development of so called machine learning methods over the last few decades (Hastie, Tibshirani, & Friedman, 2009). The subclass of methods for *supervised learning* (of which logistic regression is an example) is pertinent for classification tasks such as

actuarial risk assessment. There has been some, but not widespread, use of advanced supervised learning methods for risk assessment (Zeng et al., 2017) with some researchers finding limited benefit from the methods (Hamilton, Neuilly, Lee, & Barnoski, 2015; Tollenaar & van der Heijden, 2013). However, it is not a given that the relative performance of different supervised learning methods in a given domain can be established once and for all. The predictive performance of various methods depends on the nature of the predictors, the outcome, and the relationship between them. While more complex models often are able to extract more predictive power than more interpretable models, actual performance is unknowable before model training (Berk & Bleich, 2013; Kuhn & Johnson, 2013). To put models with different predictors on equal footing, it is thus worth considering more than a single statistical method for prediction.

Two popular methods are *elastic net* logistic regression (Zou & Hastie, 2005) and *random forest* (Breiman, 2001a). Together they represent a very small subset of available supervised learning methods but with different strengths regarding what types of the relationships between predictors and outcome they are able to detect. Trying these two methods is thus likely to capture much of the potential of a set of predictors while limiting the search for the optimal model.

Both methods have the strength that their predictive power is not diminished by non-informative predictors. This is a very suitable feature for exploring the predictive power of a set of predictors collectively as it does not require variable selection. Every method is, however, susceptible to overfitting and the use of multiple methods increases the requirement of thorough cross-validation as we describe in the method section.

We compare these methods to a more traditional method: logistic regression. Some readers will see our use of logistic regression as a straw man, arguing that logistic regression

would not be used in this manner anyway. We do not expect logistic regression to perform well when using the numerous variables in our data, but include it as a baseline measure.

Controlling for change in risk via prison term characteristics

Comparing the predictive power of static and dynamic risk factors at the beginning of the sentence is justified in a perspective that both sets of predictors reflect a snapshot at the same time point. From that time point both sets of predictors should lose some of their predictive power over time due to spontaneous, or induced, changes in risk (Andrews et al., 2011). It is, however, possible that dynamic risk factors lose their predictive power faster than static risk factors and that dynamic risk factors would show higher predictive power if updated closer to release.

In the absence of repeated measurements we can only implicitly assess the potential of updated risk assessment, but we can do this through prison term characteristics serving as proxies for change in risk. In the present study, these are decisions about conditional release and placement in open prison, supervision of parole, and crime during the sentence. These represent information that becomes available after the initial assessment and are indicators of potential changes in risk. Continuing our analysis from the perspective of predictive power, we are not interested in the nature of the relationship between these variables and recidivism (we do not consider them as risk assessment items) but rather in the added predictive information that these variables provide. Including all variables gives us an estimate of the maximum predictive power of all information available in our data. To the extent that the initial assessment fails to capture risk at release it should show up as incremental predictive validity of the prison term characteristics.

Evaluating Predictive Performance: Discrimination and Calibration

There are two parts to the predictive performance of a model: discrimination and calibration (Helmus & Babchishin, 2017). Discrimination is the ability of the test to assign higher estimated probabilities of recidivism for individuals that do in fact reoffend, compared to individuals that do not reoffend. On the other hand, a test is well calibrated if it assigns a level of estimated probabilities to a group of individuals that corresponds well with the actual recidivism rate in that group.

While good predictors are likely to improve both, discrimination and calibration are very different concepts. A test can have good discrimination without having good calibration. As long as the estimated probability is relatively higher among individuals that eventually reoffend than among individuals that do not, the discrimination can be considered good, regardless of what those estimated probabilities are. Calibration requires that the estimated probabilities correspond to actual recidivism rates. We evaluate both forms of predictive performance in the present study.

Research Questions

The present study examines the predictive power of dynamic items in the RITA and aims to contribute to the discussion of the role of dynamic predictors of recidivism in two ways. For one, we perform an explicit comparison of the predictive potential of static and dynamic predictors through methods that are designed to optimize predictive performance rather than interpretability. We believe that the methods we have chosen are very useful for quantifying the predictive power of different sets of predictors. In addition, we examine the predictive potential of dynamic variables as they are assessed in an applied setting rather than under the strict control of scale developers. This is in contrast to much of the research on dynamic risk factors (Cording

et al., 2016). We do all this both for predictions of general and of violent recidivism. Specifically, we ask the following questions:

1. Do the algorithms of elastic net and random forest result in better predictions than a more traditional method: logistic regression?
2. Do the RITA-items assessed in an applied setting in the beginning of the sentence predict recidivism?
3. Do the RITA-items have incremental predictive validity over available static predictors?
4. Does considering term characteristics as proxies for updated risk have incremental predictive validity over the assessment in the beginning of the sentence?
5. Are the models good enough to be applied?
 - a. Are the discrimination values generalizable?
 - b. Are the prediction models well calibrated?

Method

Sample

Information used in the present study was retrieved from the *National Prisoner Database* (in Finnish: *Vankitietojärjestelmä*), upheld by the Criminal Sanctions Agency in Finland. The database contains records of current and former inmates, including information about the offences they have been convicted for, temporary releases or parole, possible disciplinary reports as well as assessment reports, (if completed).

This sample consisted of a total of 1564 individuals, released from a Finnish prison between 2007 and 2011. Subjects were considered if they had the Finnish Risk and Needs Assessment Form (*Riski- ja tarvearvio*, RITA) fully completed. Of these, half ($n = 782$) were chosen based on the fact that they had been sentenced to a new prison term at the time when the

data were collected, during 2012. Each of these 782 reoffending individuals were matched by an individual of the same sex that had been released around the same date but had not been sentenced to a new term by the time the data were collected.

Between 2007 and 2011, an average of 4400 prisoners were released from Finnish prisons each year (excluding those imprisoned for remand or as substitute for unpaid fines). Out of these, 6.9% were women (Blomster, Linderborg, Muiluvuori, Salo, & Tyni, 2012). In our sample, the proportion of women was 4.3%. Because of the low number of female prisoners in our sample ($n = 68$), we decided to limit our examination to the 1496 male prisoners.

The median time to the new sentence was 102 days (max = 1321 days) and the total follow-up time had a median of 3.7 years (range = 0.5 – 5.8 years). Descriptive statistics for select predictors per reoffence category are presented in Table 1. See

<https://github.com/bennysalo/predict-recidivism> for descriptive statistics for all predictors.

Outcomes

For an individual to be considered as having reoffended this sentence had to be recorded in the prisoner database. Arguably the outcome could be called *reimprisonment* as the extraction from the database does not capture convictions with more lenient sentences. Thus, recidivism was in the present study operationalized as a new prison sentence. For these new convictions the new offence was coded according to the same categories as the current offence. A new sentence of any kind (including violent offences) was coded as *general recidivism*. New offences were considered *violent recidivism* if they included homicide or assault. We did not consider the category of the previous crime when coding for general and violent recidivism.

As a result of the sampling procedure the base rate for general recidivism was 50%. Violent recidivism, being a subset of this, had a base rate in this study of 18.6%. The follow-up

time for general recidivism is by design close to identical with the mean follow-up time for reimprisoned individuals at 3.66 years ($SD = 1.04$) compared to 3.65 ($SD = 1.05$) for individuals not reimprisoned. A quantile-quantile plot revealed very similar distributions of follow-up time for individuals reimprisoned for a violent crime ($M = 3.65$ years, $SD = 1.04$) and individuals reimprisoned for a non-violent crime or not reimprisoned at all ($M = 3.70$ years, $SD = 1.05$).

Predictors

Static predictors

Static predictors included categories for the current crime committed, prison history, and age. The offence for which the current prison was served was coded with 15 binary variables; *homicide, assault, sexual offence, offence against official authorities, other offence against the person, robbery, theft, auto theft, other property offence, criminal damage, narcotic offence, offence related to possession of weapon, traffic offence, white-collar offence, and other offences* (e.g. relating to animal maltreatment or waste mismanagement). The categories are not mutually exclusive. For example, violent resistance towards a police officer would be coded both as offence against official authorities and as assault.

Four variables pertaining to number of previous sentences were included: number of *prison terms, community service terms, remand terms, and terms as substitute for unpaid fines*. For 45 men this information was missing. For these, the number of previous sentences were coded as 0 and a new variable indicating missing information regarding previous sentences was introduced and coded as 1 for these cases.

Other variables pertaining to prison history were a binary variable indicating any *escape, unlawful absence or attempt thereof* during any of the noted terms, and the person's *age when the first term was noted in the prison database*. Current age was defined as the age at release.

Age at first term was missing for 426 individuals. Missing values were replaced by the current age (age at release) rounded down to the nearest integer (age at release was recorded as years in decimal numbers, age at first term was recorded as an integer), and like for the number of previous sentences a variable indicating missing information was introduced. Sentence length was not available in the data.

Our method for replacing missing values captures the predictive information without introducing inferred values. If the goal was to explain rather than predict recidivism other methods for handling missing data would be more appropriate. For prediction purposes the information about missingness may in itself contain implicit information that is worth retaining (Shmueli, 2010).

Dynamic predictors

The considered dynamic predictors were 50 items from the following sections of the RITA: *Accommodation and managing activities of daily living; Income and managing financial situation; Education, employment and skills supporting education and employment; Social bonds and life-style, Alcohol use; Drug use; Thinking and behavior; and Attitudes* plus the item *Takes responsibility for the current offence*. The items are all coded on a three-point scale with 0 = risk factor not present, 1 = the risk factor is somewhat present, or 2 = the risk factor is evidently present. Items can be found in the Appendix to Salo, Laaksonen, & Santtila (2016) and on <https://github.com/author/repository>.

Variables related to the prison term

The variables we used to estimate the predictive power stemming from information available after the initial assessment are *placement in open prison, granting conditional release, supervision of parole, and crime during the sentence*.

After considering the risk of absconding or crime an individual can be placed to serve part of the sentence in an open institution with minimum security. In addition to practical considerations, the individual needs to agree to abstain from drug use and, if considered necessary, take recurring drug tests. The individual can be placed back into a closed institution if required. The variable we used as predictor for recidivism was *placement at release*. Of the males, 664 concluded their sentence in an open prison. This variable is thus a predictor that is a combination of the prison personnel's view of risk, the willingness of the imprisoned individual to meet requirements, and the individual's ability to fulfill these requirements. The same is true for granting *conditional release*, where after similar considerations, the last six months of a sentence can be served in the community while under technical surveillance. For this variable, we had information available on whether conditional release was granted ($n = 337$) and whether it was eventually revoked ($n = 75$ of those granted conditional release)—this was dummy coded as two separate predictors.

In the majority of cases prisoners are released on parole. This parole is generally supervised if the remaining sentence is longer than 18 months and the supervision may take several forms. Our variable on *supervision of parole* includes minimal judgement from case workers and is primarily an indicator of sentence length. It is nevertheless related to the unfolding of the prison term and added information in relation to static and dynamic risk factors.

Finally, the variable *crime during sentence* indicates whether the individual committed or was a suspect of a crime during any lawful or unlawful absence from the prison. This also includes later reports from a new conviction where at least a part of the crime was committed during the previous sentence.

Subsets for Cross-Validation

We used two methods to validate the models, repeated cross-validation within a *training set* and validation in a separate *test set*. First, we split the sample into a training set ($n = 1122$) and a test set ($n = 374$) with three quarters of all individuals in the training set. Within the training set we used cross-validation to choose the values of so called tuning parameters (see below) and to compare the discrimination of models containing different sets of predictors. This cross-validation was done by splitting the training set into ten folds with each of the folds, in turn, serving as a validation set. This ten-fold cross-validation was repeated ten times for added accuracy of estimates of discrimination. After choosing the best values for the tuning parameters for each model (based on resulting *AUC*), the model was trained on the entire training set and then used to predict recidivism in the test set (once per model).

The generalizability of predictive performance is more conservatively evaluated in the test set. None of the observations in the test set have been involved in training or tuning the model. Validation in the smaller test set, on the other hand, has lower statistical power. Statistical power can be increased by allocating more individuals to the test set. However, as this automatically decreases the size of the training set, it diminishes the accuracy of the trained model. In prediction modelling, the balance is often struck with more individuals in the training set (Hastie et al., 2009; Kuhn & Johnson, 2013).

Machine Learning Methods

We employed two often used methods for supervised learning with differing strengths: *elastic net* logistic regression (Zou & Hastie, 2005) and *random trees* (Breiman, 2001a). The elastic net adds a penalty parameter to the estimation of a logistic regression model to reduce overfitting and thereby improve predictive power. Random forest is an aggregation of multiple

classification trees built on different bootstrapped samples and with different random selections of the predictors to consider. We compared both models to logistic regression without penalty.

Elastic net logistic regression performs well under similar conditions as logistic regression, that is, when there are linear relationships between predictors and the log-odds of the outcome that are parsimoniously captured by logistic regression coefficients. The elastic net, however, is less susceptible to the effects of multicollinearity thanks to the penalty imposed on the coefficients (Dormann et al., 2013). The strength of the random forest algorithm lies in its flexibility. By growing trees in multiple layers it can accommodate for complex interactions between predictors. However, while it splits predictors at optimal points, it does not capture linear relationships with the outcome as well as logistic regression. From an interpretability standpoint it is also inferior to logistic regression (Hastie et al., 2009; Kuhn & Johnson, 2013).

Both algorithms implement so called tuning parameters. For elastic net there are the nature and strength of the penalty parameter and for random forest there is the number of variables to randomly select for consideration in growing the classification trees. The best tuning parameter values were chosen based on cross-validation in the training set. For elastic net, we tested a broad range of possible combinations of values and then incrementally narrowed in on the best values. This procedure is documented on <https://github.com/author/repository>.

Evaluation Criteria of Predictive Performance

One of the most common metrics for evaluating discrimination is the *Area Under the Curve (AUC)* from a *Receiver Operating Characteristic (ROC)* analysis. The *AUC* considers the specificity and sensitivity of the test at all possible levels of cutoff. The *AUC* has an advantage over Pearson's r and Cohen's d in that they are not influenced by base rates or differences in within group variances (Helmus & Babchishin, 2017; Ruscio, 2008). We find Cohen's d to be a

metric that is easier to interpret when comparing models. Cohen's d is a linear metric while increasingly big differences in discrimination is needed for the same increase in AUC as it approaches 1. To maintain the robustness of AUC we do not calculate Cohen's d directly but calculate it from AUC using a formula in Table 1 of Ruscio (2008) assuming equal group sizes. This assumption does not hold for violent recidivism but makes comparison between types of recidivism easier. A table for conversions between effect sizes using this formula can also be found in Rice & Harris (2005). As a rule of thumb $d = 0.2$, 0.5 , and 0.8 are characterized as *small*, *medium* and *large* effects (Cohen, 1992). Converted to AUC by the above mentioned formula this corresponds to $AUC = 0.56$, 0.64 , and 0.71 .

Calibration statistics have unfortunately not been used as much in the risk assessment literature and the methods for examining calibration vary (Hanson, 2017; Helmus & Babchishin, 2017). Commonly calibration is evaluated in some form by comparing the expected rate of recidivism to the observed rate. One illustrative form of doing this is a calibration plot. In our calibration plot, we divided individuals into quintiles based on the estimated probability of recidivism, calculated the average estimated probability in that group, and plotted that against the de facto recidivism rate of the group. A well calibrated model follows a diagonal line where the expected and the observed recidivism rates are equal.

Statistical Software

We used the software environment *R* (Version 3.4.3; R Core Team, 2017) for all analyses. For training the model and cross-validating them in the training set we used the package *caret* (Version 6.0-78; Kuhn, 2008). Together with *caret* we used the packages *glmnet* (Version 2.0-13; Friedman, Hastie, & Tibshirani, 2010) for elastic net regression and *randomForest* (Version 4.6-12; Liaw & Wiener, 2002) for random forest. AUC values and their bootstrapped confidence

intervals were computed using the package *pROC* (Version 1.10-0; Robin et al., 2011). Plots were built using the package *ggplot2* (Version 2.2.1; Wickham, 2009).

Results

The discriminative power of the predictive models measured through *AUC* are presented in Figure 1. It illustrates the difference in the statistical power of our two validation methods. The test set validations (x-axis) have much wider and largely overlapping confidence intervals with few statistically significant differences. To compare methods and sets of predictors we rely on the more precise cross-validation figures in the training set (y-axis).

A recommended approach for development of an applied prediction tool is to choose the most promising method based on the training set and to validate the predictions in the test set (Kuhn & Johnson, 2013). Choosing the best generalizability measures based on performance in the test set would be capitalizing on chance. For each set of predictors we therefore report the best model in the training set and its performance in the the test set. This is the elastic net model in each case except for static predictors predicting violent reidivism where it is random forest. The *AUC*-values for the best models, using both validation methods are presented in Table 2.

Comparison of Methods

As expected, elastic net and random forest both produce higher discrimination than logistic regression. More importantly, the relative performance between sets of predictors is different with logistic regression than with the other two methods. A likely cause of this is that since logistic regression does not have a mechanism for handling non-informative items, it overfits more easily as the number of predictors grow. Logistic regression handles the set of 24 static predictors relatively well but performs worse with the 52 RITA-items or when the sets are

combined. The average improvement of the elastic net over an unpenalized logistic regression was equivalent to a Cohen's d of 0.18, a small effect.

For general recidivism elastic net and random forest perform almost identically. For violent recidivism prediction using static predictors seem to favor random forest and prediction using RITA-items favor elastic net. This suggests that there are interactions between static predictors that the elastic net algorithm does not capture, and that random forest is not able to fully capture the ordinal nature of the RITA-factors. The differences are small, however.

Comparison of Sets of Predictors

The RITA-items predict both general and violent recidivism at an acceptable level. The predictive power of the RITA-items is nearly identical to the predictive power of static predictors regarding prediction of general recidivism. The difference is equivalent to a Cohen's d of 0.05. When it comes to violent recidivism there is a difference equivalent to a small effect of $d = 0.25$. Combining the two sets leads to an incremental improvement in discrimination equivalent to $d = 0.13$ for general recidivism and $d = 0.04$ for violent recidivism. Finally, adding the term variables to the predictor set yields incremental discrimination equivalent to $d = 0.16$ for general recidivism and $d = 0.04$ for violent recidivism. In summation, the results suggest that the RITA-items and the term variables play a have small incremental validity over static predictors for general recidivism that this is negligible for the prediction of violent recidivism.

A helpful interpretation of AUC in the present case is the following: It is the probability that a randomly chosen re-convicted person had a higher estimated probability of re-conviction than a randomly chosen person who was not re-convicted (Hanley & McNeil, 1982; Helmus & Babchishin, 2017). The “small” difference between RITA-items and static predictors for predicting violent recidivism thus represents a 5.5% percentage point increase in this

probability—which if leveraged in impactful decisions may potentially be considered meaningful.

Generalizability

The 95% confidence intervals for discrimination in the independent test set overlaps the estimated discrimination in the training set. This adds credence to the results in the training set. These estimates are more conservative than the cross-validation estimates but the lower limit is above .7 for all chosen models except the ones using solely RITA-variables as predictors. All prediction models showed a strong relationship with recidivism with the estimated discrimination above .7.

Calibration

Figure 2 shows the calibration of the predictive model that could be considered for applied use (i.e. the models with best discrimination in the training set and excluding models using term variables). The relative performance regarding calibration corresponds roughly with the relative performance in discrimination. Elastic net and random forest models show good calibration with observed recidivism rates corresponding closely to the average expected recidivism rate in respective quintile of estimated probabilities. The differences between models are mainly in respect to the range of the estimated probabilities. The range of estimated probabilities for violent recidivism is markedly narrower with few estimations over a probability of 50%. This can partly be attributed to the lower base rate of violent recidivism. For both general and violent recidivism the range of predictions is a little narrower when using only RITA items. In both cases the more conservative range of predictions seems to be motivated. The individuals that the models are able to put in respective quintile recidivate at rates corresponding to the average estimation.

Discussion

In the introduction of the present paper we defined five questions that we aimed to answer. We first turn to discussing how we see the results answer these questions. After that we discuss what the results suggest about the role of dynamic variables and to policy implications for predicting recidivism.

Performance of Elastic Net and Random Forest

Elastic net and random forest both produced considerably better predictions than logistic regression. This is expected since a traditional logistic regression is susceptible to overfitting (Hastie et al., 2009). When examining the predictive power of different sets of predictors the most important drawback is that different sets of predictors are affected differently by the shortcomings of logistic regression. The difference in predictive power of the different models does not only depend on the predictive information contained in the predictors but also on their suitability for multiple logistic regression.

There are other ways to use logistic regression that could be considered that would arguably be better. For example, we could have attempted to select a subset of uncorrelated items or used a dimension reduction technique to create fewer predictors that would have been more suitable for use in logistic regression (Dormann et al., 2013; Hastie et al., 2009). This would, however, have added an extra step in our analyses that might have introduced subjective decisions, loss of information, or both.

We would expect the random forest model to outperform the elastic net when the relationship between the predictors used and the outcome is complex, involving multiple layers of interactions. In our case, the elastic net algorithm produced the highest level of discrimination in all cases except one: the random forest model produced better discrimination when using only

static items to predict violent recidivism. An interpretation of this is that multiple items make a small univariate contribution and are best used simply added to each other (which an elastic net effectively does). Only when limiting the available information to the static items does the ability of random forest to account for complex interactions make a difference.

Predictive Power of RITA-items

Estimated probabilities based on RITA-items showed a strong relationship with both general and violent recidivism. The discrimination is on par with instruments with high discrimination in their validation studies (see Table 1 in Cording et al., 2016) and above the average reported in meta-analyses (Campbell et al., 2009; Gendreau, Little, & Goggin, 1996; Katsiyannis, Whitford, Zhang, & Gage, 2018; Olver, Stockdale, et al., 2014; Walters, 2003; Yang et al., 2010). This is notable since the RITA is administered with considerable freedom of interpretation of the case worker and has not gone through the process of thorough construct validation that many other instruments have. These results suggest that the level of discrimination that is seen in research settings can be reached by dynamic risk factors also in an applied setting.

We attribute the relatively high discrimination to two factors working in tandem: customized weighing of items and thorough cross-validation. We did not use a predefined scoring scheme for summing up the items scores but instead let the algorithms find the best solution, and separately for violent and general recidivism. This means that the items can be used in a way that best links the items as interpreted by the case workers to recidivism in this Finnish sample. The elastic net and random forest algorithms are both well suited for this task. However, we want the weighting to be customized to this particular setting but not to this particular sample. Cross-validation helps avoid this overfitting. The models were tuned in the training set by repeatedly

testing discrimination in a subsample that was not used to train the model. This means that when the models were tested on entirely new data (the test set) they performed at the level expected by the results from the training data.

Incremental Predictive Power of RITA-Items Over Static Predictors

The same factors that benefit the predictive power of the RITA-items also benefit the static predictors and they performed on the same level and better than RITA-items. Adding RITA-items to the static predictors gave small improvement in discrimination for general recidivism, but not negligible for violent recidivism.

One interpretation of this is that using static predictors alone comes close to filling the predictive potential of the information available at the beginning of the sentence. The difference is very small when it comes to predicting violent recidivism and dynamic items therefore have little to add. This interpretation fits well the idea of Coid et al. (2011) about a 'glass ceiling'. The glass ceiling in the present study, however, lies around $AUC = .8$ in comparison to the ceiling of $.72$ for prospective studies of violent recidivism that Coid et al. identified.

Incremental Predictive Power of Term Variables

Adding variables related to the unfolding of the prison term breaks the aforementioned glass ceiling with a small improvement of the prediction of general recidivism. Interestingly, we do not see the same with violent recidivism. It is worth noting that what is measured here is discrimination. A lack of increase in discriminative power should not be interpreted as meaning that what happens, or what is learned, during the sentence does not have an effect on the risk of violence. No change in discrimination means no change in the ranking according to estimated risk of recidivism. Thus, a better interpretation, than a claim of no effect of treatment, might be that extra care is put in the supervision decisions regarding violent offenders and that high risk

individuals are treated in line with readily available information about risk. For example, even if placement in open prison would be a protective factor against recidivism, adding placement as a predictor in the prediction model does not affect the discriminative power if placement in open prison is primarily granted to individuals with low risk of recidivism in line with the actuarial assessment.

To be clear, the present study makes no claim of being able to assess the efficacy of any specific intervention or prison treatment in general. How the prison term unfolds may have important consequences. The high discriminative power of static items simply implies that those who start off with the highest risk, remain the ones with the highest risk. This will be true even if low risk individuals have their risk lowered further, or if the risk of recidivism is lowered collectively for the prison population.

Generalizability and Calibration of the Models

Validation of the chosen models, in an independent test set, shows that discrimination is strong enough for the models to be applied, especially any model that includes static predictors. Effect sizes should be interpreted in context, however. The interpretation of, for example, $AUC = .78$ is that, in 78% of cases, a randomly picked individual that is reimprisoned has a higher estimated probability of reoffending than a randomly picked individual that is not reimprisoned. This is a helpful, but not a perfect measure.

While AUC is a metric robust to differences in base rate, the prediction of violent recidivism is harder simply because there is less information about violent reoffenders. In this study the 278 individuals with new violent crimes provided enough information to reach a level of discrimination for violent recidivism on par with discrimination for general recidivism.

Base rate differences influenced calibration more. Calibration was good for all the models but the range of estimated probabilities was more limited for violent recidivism. The models were not able to identify any individual with a probability of committing a new violent crime much higher than 50%. This may also correspond to the true predictability of violent offences. While estimated probabilities tend to be higher for individuals committing new violent crimes, these estimations are lower overall because of a lower base rate. Using only RITA items also offered a narrower range of estimated probabilities for both general and violent recidivism making models that include static items, alone or in combination with dynamic items, a little more useful.

The Role of Dynamic Predictors

The patterns of discriminative power of the models tell different stories about the predictability of general and violent recidivism. Dynamic items and how the sentence unfolds seems to be able to tell us more about the risk for general recidivism than about the risk for violent recidivism. For violent recidivism we seem to be able to learn much of what there is to learn, simply from the static variables. There is also a bigger incremental gain from adding term related predictors to the predictions for general recidivism than violent recidivism.

This suggests that general recidivism is more driven by changable circumstances than violent recidivism is. While circumstances can be stubbornly stable in themselves, the stable personal characteristics that have resulted in previous criminal behavior (and other static markers) are likely to have a relatively bigger role when explaining possible future violence. Put another way, risk for general recidivism appears to be somewhat more dynamic than risk for violent recidivism.

This study is limited in what it can say about the potential of updating assessment of risk during a prison sentence using RITA-items. Our four term variables provide an imperfect measure of changes in risk and there is no substitute for repeated measures. What remains clear however is that there is, regardless, a high value of being able to update risk assessments.

Policy Implications

The static predictors in the present study performed a little better in relation to the dynamic predictors than in meta-analyses on prediction on general (Gendreau et al., 1996) and violent recidivism (Campbell et al., 2009). The results are, however, in line with previous conclusions that show that dynamic and static risk are nearly interchangeable when it comes to the limited task of predicting reoffending (Campbell et al., 2009; Gendreau et al., 1996; Yang et al., 2010). There has been several calls for a stronger focus on development of risk assessment tools to allow causal inferences about recidivism (Cording et al., 2016; Monahan & Skeem, 2016; Ward, 2016). For this the measured construct need to be truly dynamic. That changes in risk can be measured and that these changes have incremental predictive power has been demonstrated using the *Violence Risk Scale* (Lewis et al., 2013; Olver, Beggs Christofferson, & Wong, 2015) *Historical Clinical Risk Mangement-20* and *Structured Assessment of Protective Factors for Violence Risk* (De Vries Robbé et al., 2015).

The predictive power of the RITA-items suggest some potential of developing the RITA in a similar manner. At the same time, the relatively strong predictive power of the static predictors suggest that predictive power does not have to be a primary concern when developing the RITA. Development of the RITA can focus on establishing strong construct validity which strenghtens its usability for its current use of identifying treatment targets. If improved construct validity and reliability opens up for measurements of changes in risk, those might provide

stronger incremental predictive power. However, the initial assessment can rely primarily on static predictors. If RITA is developed further, new items can be retested in relationship to the static predictors. Customization of weights might be worthwhile allowing different scoring rules when identifying treatment targets and when assessing risk. In other words, potential development of RITA should focus on ensuring that the measures are fair, informative and reliable. Increased predictive power might be a side effect of this, but does not need to be for the purposes of an initial risk assessment.

Predictions using the current predictors can be made using a computer algorithm or converted into an analog tool. Because elastic net performed very well in this case, and produces coefficients that are interpretable in the same way as in a traditional logistic regression, an analog tool can easily be developed. This conversion might entail simplifying coefficients, deleting low information items and possibly complementing static items with the best dynamic items to reach a parsimonious model that performs close to the best models including all static and dynamic items and using computerized algorithms.

Limitations and Further Research

It is worth reminding ourselves that predictions work better on a group level than on an individual level. The calibration plots show that predictions are quite accurate averaged over the members of a certain risk group. The model can allocate an individual to the wrong risk category without severely damaging the performance on a group level.

This study examines only the static and dynamic variables available to us. Other static items, and especially other dynamic items might have given added predictive power. The strength of this study is however in its applied setting.

Even more so, the items that aim to capture what happens during the sentence was limited to only four variables. With repeated assessment of dynamic items more could be said about how dynamic risk factors change and how the risk of recidivism change. With the information already at hand we could go beyond discriminative power and also investigate the effect of the practices of placement in open prison and conditional release on the risk of recidivism. By matching groups on their propensity for being placed in open prison or granted conditionals release we could examine if the level of supervision change the level of risk within a risk group.

References

- Ægisdóttir, S., White, M. J., Spengler, P. M., Maugherman, A. S., Anderson, L. A., Cook, R. S., ... Rush, J. D. (2006). The meta-analysis of clinical judgment project: Fifty-six years of accumulated research on clinical versus statistical prediction. *The Counseling Psychologist*, 34(3), 341–382. doi: 10.1177/0011000005285875
- Andrews, D. A., Bonta, J., & Wormith, J. (2004). *Manual for the Level of Service/Case Management Inventory (LS/CMI)*. Toronto: Multi-Health Systems.
- Andrews, D. A., Bonta, J., & Wormith, J. S. (2006). The recent past and near future of risk and/or need assessment. *Crime & Delinquency*, 52(1), 7–27. doi: 10.1177/0011128705281756
- Andrews, D. A., Bonta, J., Wormith, J. S., Guzzo, L., Brews, A., Rettinger, J., & Rowe, R. (2011). Sources of variability in estimates of predictive validity: A specification with Level of Service general risk and need. *Criminal Justice and Behavior*, 38(5), 413–432. doi: 10.1177/0093854811401990
- Andrews, D. A., Zinger, I., Hoge, R. D., Bonta, J., Gendreau, P., & Cullen, F. T. (1990). Does correctional treatment work? A clinically relevant and psychologically informed meta-analysis. *Criminology*, 28(3), 369–404. doi: 10.1111/j.1745-9125.1990.tb01330.x
- Berk, R. A., & Bleich, J. (2013). Statistical procedures for forecasting criminal behavior. *Criminology & Public Policy*, 12(3), 513–544. doi: 10.1111/1745-9133.12047
- Blomster, P., Linderborg, H., Muiluvuori, M.-L., Salo, I., & Tyni, S. (2012). *Rikosseuraamuslaitoksen tilastoja 2011 [Statistics of the Criminal Sanctions Agency 2011]*. Helsinki: Rikosseuraamuslaitos.
- Breiman, L. (2001a). Random Forests. *Machine Learning*, 45(1), 5–32. doi:

10.1023/A:1010933404324

Breiman, L. (2001b). Statistical modeling: The two cultures. *Statistical Science*, 16(3), 199–231.

doi: 10.1214/ss/1009213726

Campbell, M. A., French, S., & Gendreau, P. (2009). The prediction of violence in adult offenders: A meta-analytic comparison of instruments and methods of assessment. *Criminal Justice and Behavior*, 36(6), 567–590. doi: 10.1177/0093854809333610

Clarke, M. C., Peterson-Badali, M., & Skilling, T. A. (2017). The relationship between changes in dynamic risk factors and the predictive validity of risk assessments among youth offenders. *Criminal Justice and Behavior*, 44(10), 1340–1355. doi: 10.1177/0093854817719915

Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159. doi: 10.1037/0033-2909.112.1.155

Coid, J. W., Yang, M., Ullrich, S., Zhang, T., Sizmur, S., Farrington, D., & Rogers, R. (2011). Most items in structured risk assessment instruments do not predict violence. *Journal of Forensic Psychiatry & Psychology*, 22(1), 3–21. doi: 10.1080/14789949.2010.495990

Cording, J. R., Beggs Christofferson, S. M., & Grace, R. C. (2016). Challenges for the theory and application of dynamic risk factors. *Psychology, Crime & Law*, 22(1–2), 84–103. doi: 10.1080/1068316X.2015.1111367

Criminal Sanctions Agency in Finland. (2004). *Sijoittajaysikkötoiminnan käsikirja. Rikosseuraamusalan käsikirjoja 1/2004 [Operational handbook of the assessment and allocation unit. Handbooks of the criminal sanctions field 1/2004]*.

Dawes, R., Faust, D., & Meehl, P. (1989). Clinical versus actuarial judgment. *Science*, 243(4899), 1668–1674. doi: 10.1126/science.2648573

- De Vries Robbé, M., de Vogel, V., Douglas, K. S., & Nijman, H. L. I. (2015). Changes in dynamic risk and protective factors for violence during inpatient forensic psychiatric treatment: Predicting reductions in postdischarge community recidivism. *Law and Human Behavior*, 39(1), 53–61. doi: 10.1037/lhb0000089
- Debidin, M. (Ed.). (2009). *A compendium of research and analysis on the Offender Assessment System (OASys) 2006-2009*. Ministry of Justice, United Kingdom. Retrieved from <http://webarchive.nationalarchives.gov.uk/20110201125714/http://www.justice.gov.uk/publications/docs/research-analysis-offender-assessment-system.pdf>
- Dormann, C. F., Elith, J., Bacher, S., Buchmann, C., Carl, G., Carré, G., ... Lautenbach, S. (2013). Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. *Ecography*, 36(1), 27–46. doi: 10.1111/j.1600-0587.2012.07348.x
- Flores, A. W., Lowenkamp, C. T., Holsinger, A. M., & Latessa, E. J. (2006). Predicting outcome with the Level of Service Inventory-Revised: The importance of implementation integrity. *Journal of Criminal Justice*, 34(5), 523–529. doi: 10.1016/j.jcrimjus.2006.09.007
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1–22.
- Gendreau, P., Little, T., & Goggin, C. (1996). A meta-analysis of the predictors of adult offender recidivism: What works? *Criminology*, 34(4), 575–608. doi: 10.1111/j.1745-9125.1996.tb01220.x
- Gottfredson, S. D., & Moriarty, L. J. (2006). Statistical risk assessment: Old problems and new applications. *Crime & Delinquency*, 52(1), 178–200. doi: 10.1177/0011128705281748
- Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: a meta-analysis. *Psychological Assessment*, 12(1), 19–30. doi:

10.1037/1040-3590.12.1.19

Hamilton, Z., Neuilly, M.-A., Lee, S., & Barnoski, R. (2015). Isolating modeling effects in offender risk assessment. *Journal of Experimental Criminology*, *11*(2), 299–318. doi:

10.1007/s11292-014-9221-8

Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, *143*(1), 29–36. doi:

10.1148/radiology.143.1.7063747

Hanson, R. K. (2017). Assessing the calibration of actuarial risk scales: A primer on the E/O index. *Criminal Justice and Behavior*, *44*(1), 26–39. doi: 10.1177/0093854816683956

Hanson, R. K., & Harris, A. (2000). *The Sex Offender Need Assessment Rating (SONAR): A method for measuring change in risk levels 2000-1*. Public Works and Government Services Canada. Retrieved from <https://www.publicsafety.gc.ca/cnt/rsrscs/pblctns/sx-ffndr-nd/index-en.aspx>

Hare, R. D. (2003). *The Hare Psychopathy Checklist - Revised*. (2nd ed.). Toronto: Multi-Health Systems.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning*. New York, NY: Springer New York. doi: 10.1007/b94608

Helmus, L. M., & Babchishin, K. M. (2017). Primer on risk assessment and the statistics used to evaluate its accuracy. *Criminal Justice and Behavior*, *44*(1), 8–25. doi:

10.1177/0093854816678898

Howard, P. D., & Dixon, L. (2013). Identifying change in the likelihood of violent recidivism: causal dynamic risk factors in the OASys violence predictor. *Law and Human Behavior*, *37*(3), 163–174. doi: 10.1037/lhb0000012

- Jones, N. J., Brown, S. L., & Zamble, E. (2010). Predicting criminal recidivism in adult male offenders: Researcher versus parole officer assessment of dynamic risk. *Criminal Justice and Behavior*, 37(8), 860–882. doi: 10.1177/0093854810368924
- Katsiyannis, A., Whitford, D. K., Zhang, D., & Gage, N. A. (2018). Adult recidivism in United States: A meta-analysis 1994–2015. *Journal of Child and Family Studies*, 27(3), 686–696. doi: 10.1007/s10826-017-0945-8
- Kraemer, H. C. (2003). Current concepts of risk in psychiatric disorders. *Current Opinion in Psychiatry*, 16(4), 421–430. doi: 10.1097/01.yco.0000079210.36371.52
- Kroner, D. G., Mills, J. F., & Reddon, J. R. (2005). A Coffee Can, factor analysis, and prediction of antisocial behavior: The structure of criminal risk. *International Journal of Law and Psychiatry*, 28(4), 360–374. doi: 10.1016/j.ijlp.2004.01.011
- Kuhn, M. (2008). Building predictive models in R using the caret package. *Journal of Statistical Software*, 28(5). doi: 10.18637/jss.v028.i05
- Kuhn, M., & Johnson, K. (2013). *Applied predictive modeling*. New York, NY: Springer New York. doi: 10.1007/978-1-4614-6849-3
- Lewis, K., Olver, M. E., & Wong, S. C. P. (2013). The Violence Risk Scale. *Assessment*, 20(2), 150–164. doi: 10.1177/1073191112441242
- Liaw, A., & Wiener, M. (2002). Classification and regression by randomForest. *R News*, 2(3), 18–22. Retrieved from <http://cran.r-project.org/doc/Rnews/>
- Lilja, T. (2014). *Valvotun koevapauden peruuntumisen ennustettavuus [The predictability of annulment of conditional release]*. Laurea University of Applied Sciences, Tikkurila, Finland. Retrieved from https://www.theseus.fi/bitstream/handle/10024/85732/Lilja_Tiina.pdf?sequence=1

- Monahan, J., & Skeem, J. L. (2016). Risk assessment in criminal sentencing. *Annual Review of Clinical Psychology*, 12(1), 489–513. doi: 10.1146/annurev-clinpsy-021815-092945
- Nuffield, J. (1982). *Parole decision-making in Canada: Research towards decision guidelines*. Ottawa: Solicitor General of Canada.
- Olver, M. E., Beggs Christofferson, S. M., Grace, R. C., & Wong, S. C. P. (2014). Incorporating change information into sexual offender risk assessments using the Violence Risk Scale–Sexual Offender Version. *Sexual Abuse: A Journal of Research and Treatment*, 26(5), 472–499. doi: 10.1177/1079063213502679
- Olver, M. E., Beggs Christofferson, S. M., & Wong, S. C. P. (2015). Evaluation and applications of the clinically significant change method with the Violence Risk Scale–Sexual Offender version: implications for risk-change communication. *Behavioral Sciences & the Law*, 33(1), 92–110. doi: 10.1002/bsl.2159
- Olver, M. E., Stockdale, K. C., & Wormith, J. S. (2014). Thirty years of research on the Level of Service Scales: A meta-analytic examination of predictive accuracy and sources of variability. *Psychological Assessment*, 26(1), 156–176. doi: 10.1037/a0035080
- Quinsey, V. L. (2009). Are we there yet? Stasis and progress in forensic psychology. *Canadian Psychology/Psychologie Canadienne*, 50(1), 15–21. doi: 10.1037/a0014401
- Quinsey, V. L., Harris, G. T., Rice, M. E., & Cormier, C. A. (2006). *Violent offenders: Appraising and managing risk* (2nd ed.). Washington, DC: American Psychological Association.
- R Core Team. (2017). R: A language and environment for statistical computing. *R Foundation for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Retrieved from <http://www.r-project.org>

- Rice, M. E., & Harris, G. T. (2005). Comparing effect sizes in follow-up studies: ROC Area, Cohen's d , and r . *Law and Human Behavior*, 29(5), 615–620. doi: 10.1007/s10979-005-6832-7
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & Müller, M. (2011). pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*, 12(1), 77. doi: 10.1186/1471-2105-12-77
- Ruscio, J. (2008). A probability-based measure of effect size: Robustness to base rates and other factors. *Psychological Methods*, 13(1), 19–30. doi: 10.1037/1082-989X.13.1.19
- Salo, B., Laaksonen, T., & Santtila, P. (2016). Construct validity and internal reliability of the Finnish risk and needs assessment form. *Journal of Scandinavian Studies in Criminology and Crime Prevention*, 17(1). doi: 10.1080/14043858.2016.1161940
- Shmueli, G. (2010). To explain or to predict? *Statistical Science*, 25(3), 289–310. doi: 10.1214/10-STS330
- Singh, J. P., Grann, M., & Fazel, S. (2013). Authorship bias in violence risk assessment? A systematic review and meta-analysis. *PloS One*, 8(9), e72484. doi: 10.1371/journal.pone.0072484
- Tollenaar, N., & van der Heijden, P. G. M. (2013). Which method predicts recidivism best?: A comparison of statistical, machine learning and data mining predictive models. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 176(2), 565–584. doi: 10.1111/j.1467-985X.2012.01056.x
- Vankeuslaki [Impisonment Act], 767 (2005). Finland.
- Walters, G. D. (2003). Predicting criminal justice outcomes with the Psychopathy Checklist and Lifestyle Criminality Screening Form: a meta-analytic comparison. *Behavioral Sciences &*

- the Law*, 21(1), 89–102. doi: 10.1002/bsl.519
- Ward, T. (2016). Dynamic risk factors: scientific kinds or predictive constructs. *Psychology, Crime & Law*, 22(1–2), 2–16. doi: 10.1080/1068316X.2015.1109094
- Ward, T., & Beech, A. R. (2015). Dynamic risk factors: a theoretical dead-end? *Psychology, Crime & Law*, 21(2), 100–113. doi: 10.1080/1068316X.2014.917854
- Webster, C. D., Douglas, K. S., Eaves, D., & Hart, S. D. (1997). *HCR-20: Assessing risk of violence, version 2*. Vancouver: Mental Health Law and Policy Institute, Simon Fraser University.
- Wickham, H. (2009). *ggplot2: Elegant graphics for data analysis*. New York, NY: Springer. doi: 10.1007/978-0-387-98141-3
- Yang, M., Wong, S. C. P., & Coid, J. W. (2010). The efficacy of violence prediction: a meta-analytic comparison of nine risk assessment tools. *Psychological Bulletin*, 136(5), 740–767. doi: 10.1037/a0020473
- Zeng, J., Ustun, B., & Rudin, C. (2017). Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(3), 689–722. doi: 10.1111/rssa.12227
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. doi: 10.1111/j.1467-9868.2005.00503.x

Table 1

Mean (and standard deviation) or frequency (and percentage) for select predictors by reoffense category

	No reoffence (n = 748)	Non-violent reoffence (n = 470)	Violent reoffence (n = 278)
Age at first term ^a	29.5 (11.5)	22.8 (6.3)	22.4 (6.4)
Age at release	38.3 (11.0)	33.6 (9.1)	32.2 (9.4)
Previous prison terms ^b	2.1 (3.3)	4.0 (4.2)	3.7 (4.5)
Current offence ^c			
Homicide	67 (9%)	10 (2.1%)	17 (6.1%)
Assault	413 (55.2%)	309 (65.7%)	201 (72.3%)
Sexual offence	91 (12.2%)	8 (1.7%)	10 (3.6%)
Theft	137 (18.3%)	160 (34%)	142 (51.1%)
Narcotic offence	167 (22.3%)	136 (28.9%)	89 (32%)
Placement in open prison	426 (57%)	169 (36%)	69 (24.8%)
Conditional release			
No conditional release	511 (68.3%)	407 (86.6%)	241 (86.7%)
Successful	197 (26.3%)	43 (9.1%)	22 (7.9%)
Cancelled	40 (5.3%)	20 (4.3%)	15 (5.4%)
Parole supervised	525 (70.2%)	310 (66%)	175 (62.9%)

Note. ^aMissing values for 426 individuals excluded. ^bMissing values for 45 individuals excluded. ^cCategories are not mutually exclusive. For descriptive statistics for all predictors see <https://github.com/bennysalo/predict-recidivism>.

Table 2

Discrimination of the best models in training set per outcome and predictor set

Predictor set	Cross-validation in training set				Validation in test set			
	AUC			<i>d</i>	AUC			<i>d</i>
	Mean	LL	UL		Estimate	LL	UL	
General recidivism								
RITA-items	.760	.751	.770	1.00	.725	.676	.778	0.85
Static	.772	.763	.780	1.05	.785	.738	.829	1.12
Static and RITA-items	.799	.790	.808	1.19	.782	.734	.827	1.10
All incl. term variables	.829	.821	.837	1.34	.804	.760	.847	1.21
Violent recidivism								
RITA-items	.736	.723	.748	0.89	.740	.677	.797	0.91
Static ¹	.791	.782	.800	1.15	.798	.737	.851	1.18
Static and RITA-items	.799	.788	.809	1.18	.803	.744	.855	1.21
All incl. term variables	.806	.796	.816	1.22	.806	.752	.857	1.22

Note. ¹Best model is random forest. All other models are elastic net. AUC = Area Under Curve, LL = lower limit, UL = upper limit of 95% confidence interval. *d* = Cohen's *d* calculated from AUC using formula from Table 1 in Ruscio (2008) assuming equal group sizes. This assumption does not hold for violent recidivism but makes comparison between types of recidivism easier.

Figure 1. Discrimination of the predictive models measured in Area Under the Curve (AUC).

AUC estimated in the test set on the is plotted on the x-axis against AUC estimated through 10-fold cross-validation repeated 10 times in the training set on the y-axis. Lines around points represent 95% confidence interval of the respective estimate. The diagonal dashed line represents equal performance evaluated via both methods of validation. Note that the range of both axes are truncated to make the small differences visible.

Figure 2. Calibration plot for the predictive models. The test sample ($n = 374$) is divided into quintiles based on the estimated probability of recidivism. The average estimated probability in respective quintile is plotted on the x-axis against the de facto recidivism rate in the quintile on the y-axis. The diagonal line represents perfect calibration where estimated and de facto recidivism rates are equal. Note that individuals may belong to different quintiles depending on model.