

Is AI different for SE?

Amritanshu Agrawal, Tim Menzies
aagrawa8@ncsu.edu, timm@ieee.org
North Carolina State University, USA

ABSTRACT

What AI tools are needed for SE? Ideally, we should have simple rules that peek at data, then say “use this tool” or “use that tool”.

To find such a rule, we explored 120 different data sets addressing numerous problems, including bad smell detection, predicting Github issue close time, bug report analysis, defect prediction and dozens of other non-SE problems. To this data, we apply a SE-based tool that (a) out-performs the state-of-the-art for these SE problems yet (b) fails very badly on standard AI problems.

In those results, we can find a simple rule for when to use/avoid the SE-based tool. SE data is often about infrequent issues, like the occasional defect, or the rarely exploited security violation, or the requirement that holds for one special case. But as we show, standard AI tools work best when the target is relatively more frequent. Also, we can exploit this special properties of SE, to great effect (to rapidly find better optimizations for SE tasks via a tactic called “dodging”, explained in this paper).

More generally, this result says we need a new kind of SE research for developing new AI tools that are more suited to SE problems.

1 INTRODUCTION

The last 15 years have seen a dramatic increase in the application of AI tools on software engineering— as evidenced by numerous highly cited publications [30, 38, 43, 47, 48, 51, 54, 69, 75] and the growing prominence of venues such as the Mining Software Repositories conference (now ranked in the top dozen software systems venue by Google Scholar). Researchers reported that these AI tools are often used in a “black box” manner without reflecting on the merits of choices associated with a particular tool [9, 23].

Such black-box usage is risky since it means SE practitioners might be applying the wrong tool. An emerging question in the application of AI techniques to software engineering is the validity of the assumptions that underlie the creation of those tools and techniques [3, 9, 23, 58, 71]. These articles show superior results can be obtained by specializing AI tools to the particulars of software engineering problem (see examples in next section).

One way to specialize AI tools for software engineering is *hyperparameter optimization* that automatically search through the space of possible learner settings to discover, for example, how many clusters work best for K-means. Table 1 shows some of the hyperparameter options seen in recent SE papers. The search space is very large. Assuming that the numerics of Table 1 are divided into ten bins, then there are billions of different combinations in that table. Hence, these optimizers can be very slow to execute.

Also, even when they work, these optimizers **do not offer any generalization across their results**. Hyperparameter optimization does not offer general rules for when some AI tools are more useful than others. This is significant deficiency. If such an insight

was available, we could offer engineers a simple rule that lets them quickly select appropriate AI tools.

To address this deficiency, this paper applies various tools to numerous problems from SE and elsewhere. Some of these tools are general AI tools while another, called DODGE [1], was developed by Agrawal et al. specifically for SE applications. This tool is a hyperparameter optimizer that auto-adjusts learner settings in a somewhat unusual way (see the “relax” heuristics of §3). Agrawal et al. [1] justified that approach via certain SE-based properties and some limited experimentation of 16 data sets from two SE problems (bug report analysis and defect prediction).

Agrawal et al. did not test their supposed “SE-based” tool on non-SE problems. Therefore we apply DODGE to 120 data sets from SE and elsewhere. Specifically, 37 non-SE problems from the standard AI literature (taken from a standard machine learning data repository) as well as 83 SE problems which includes 6 issue tracking data sets, 10 defect prediction data sets 63 data sets exploring Github issue close time; and 4 data sets exploring bad smell detection.

We see that DODGE **fails badly** on non-SE problems but **performs very well** on SE problems. From that experience, we can:

- Offer a method for deciding when to use/avoid general AI tools.
- Describe a case study utilizing that methodology.
- Show that this method can indeed recommend when to use a general AI tool or a SE-based tool.
- Show that this recommendation is useful and if we follow them, we can find better optimizations for SE problems.
- Show that SE data is very different to the data used to develop many AI tools. SE problem data is often about infrequent issues, for e.g., the occasional defect, the rarely exploited security violation, the requirement that only holds for one small part of a system. But as shown by our results, many standard AI tools work best when the target is relatively more frequent.

The rest of this paper is structured as follows. The next section offers notes on related work and some standard AI tools. This is followed by a description of DODGE in §3. After that, in §4, we describe the 120 data sets used to test DODGE as well as other more standard AI tools. The experiments of §5 illustrate how well DODGE performs for SE problems and how badly it performs otherwise. The meta-analysis of §6 generates the rule that explains why DODGE works so well/poorly for SE/other problem types.

Our conclusion is that the algorithms which we call “general AI tools” may not be “general” at all. Rather, they are tools which are powerful in their home domain but need to be used with care if applied to some new domains like software engineering. Hence, we argue that it is not good enough to just take AI tools developed elsewhere, then apply them verbatim to SE problems. Software engineers need to develop AI tools that are better suited to the particulars of SE problems.

As to the relationship to prior work, Table 1 and §3 are shortened version of material from Agrawal et al. [1]. Also, of the 120 data sets

processed here, 16 come from that prior study— see the results in §5.4.1 and §5.4.2 (but just as a quality assurance measure, we reran all the experiments for those 16 data sets). The rest of this paper is new, including the experiments on 104 new data sets, and the meta-analysis of §6.

2 BACKGROUND

2.1 Lessons from Applying Standard AI Tools

This section offers examples where other SE researchers have argued that software engineering is different (so we need our own kind of specialized AI tools).

Binkley et al. [9] note that information retrieval tools for SE often equates word frequency with word importance, even though the number of occurrences of a variable name such as “tmp” is not necessarily indicative of its importance. They argue that the negative impacts of such differences manifest themselves when off-the-shelf IR tools are applied in the software domain.

Software analytics usually applies standard AI data mining algorithms as “black boxes” [9, 23], where researchers do not tinker with internal design choices of those AI tools. This is not ideal. Much recent research has shown that such automatic hyperparameter optimizers can automate that tinkering process, and greatly improve predictive performances [3, 4, 23, 61, 71].

Another example comes from sentiment analysis. Standard sentiment analysis tools are usually trained on non-SE data (e.g., the Wall Street Journal or Wikipedia). Novielli et al. [58] recently developed their own sentiment analysis for the software engineering domain. After re-training those tools on an SE corpus, they found not only better performance at predicting sentiment, but also more agreement between different sentiment analysis tools.

Yet another example of “SE needs different AI tools” comes from the “naturalness” work of Devanbu et al [32]. Software has the property that some terms are used very often and everything else is used with exponentially less frequency. This means that n-gram “language models” (where frequency counts of the last N terms are used to predict the next term) can be very insightful for (a) recommending what should be used next; or (b) predicting that the next token is unusual and possibly erroneous [67]. Note that this “natural” approach can be more insightful than standard AI tools such as deep learning [31].

While the above work certainly illustrates the main point of this paper, it does not offer a simple general rule to look at new data and decide when to use pre-existing standard AI tools or when to move on to other technologies. This paper tried answering the same.

2.2 Standard AI Tools

This section describes some standard AI tools: grid search, DE, GA, SVM, Random Forests, decision tree learners, logistic regression, Naive Bayes, LDA, and EM. All these algorithms are examples of general AI tools developed outside of software engineering and then applied to SE. That is, one way to view SE is that it is just a problem domain where general AI tools can be deployed.

This paper takes another view. We say software engineering is different and that those differences mean that best results arise from tools specialized to the particulars of software engineering. One such tool is DODGE (described in the next section).

Table 1: Hyperparameter options seen in recent SE papers [2, 3, 23, 25] and in the documentation of a widely-used data mining library (Scikit-learn [62]) for a list of options not explored in those studies. While this list is incomplete, it does include many of the hyperparameter optimizations being explored in the literature.

<p>Learners:</p> <ul style="list-style-type: none"> DecisionTreeClassifier(criterion=b, splitter=c, min_samples_split=a) <ul style="list-style-type: none"> a, b, c = randuniform(0.0,1.0), randchoice(['gini','entropy']), randchoice(['best','random']) RandomForestClassifier(n_estimators=a,criterion=b, min_samples_split=c) <ul style="list-style-type: none"> a,b,c = randint(50, 150), randchoice(['gini', 'entropy']), randuniform(0.0, 1.0) LogisticRegression(penalty=a, tol=b, C=float(c)) <ul style="list-style-type: none"> a,b,c=randchoice(['l1','l2']), randuniform(0.0,0.1), randint(1,500) MultinomialNB(alpha=a) = randuniform(0.0,0.1) KNeighborsClassifier(n_neighbors=a, weights=b, p=d, metric=c) <ul style="list-style-type: none"> a, b, c = randint(2, 25), randchoice(['uniform', 'distance']), randchoice(['minkowski','chebyshev']) if c='minkowski': d= randint(1,15) else: d=2 	175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199
<p>Pre-processors for defect prediction, Issue lifetime, Bad Smells:</p> <ul style="list-style-type: none"> StandardScaler MinMaxScaler MaxAbsScaler RobustScaler(quantile_range=(a, b)) = randint(0,50), randint(51,100) KernelCenterer QuantileTransformer(n_quantiles=a, output_distribution=c, subsample=b) <ul style="list-style-type: none"> a, b = randint(100, 1000), randint(1000, 1e5) c = randchoice(['normal','uniform']) Normalizer(norm=a) = randchoice(['l1', 'l2','max']) Binarizer(threshold=a) = randuniform(0,100) SMOTE(a=n_neighbors, b=n_synthetic, c=Minkowski_exponent) <ul style="list-style-type: none"> a, b = randint(1,20),randchoice(50,100,200,400) c = randuniform(0.1,5) 	200 201 202 203 204 205 206 207 208 209 210 211 212 213 214
<p>Pre-Processors for Text mining:</p> <ul style="list-style-type: none"> CountVectorizer(max_df=a, min_df=b) = randint(100, 1000), randint(1, 10) TfidfVectorizer(max_df=a, min_df=b, norm=c) <ul style="list-style-type: none"> a, b, c = randint(100, 1000), randint(1, 10), randchoice(['l1', 'l2', None]) HashingVectorizer(n_features=a, norm=b) <ul style="list-style-type: none"> a = randchoice([1000, 2000, 4000, 6000, 8000, 10000]) b = randchoice(['l1', 'l2', None]) LatentDirichletAllocation(n_components=a, doc_topic_prior=b, topic_word_prior=c, learning_decay=d, learning_offset=e, batch_size=f) <ul style="list-style-type: none"> a, b, c = randint(10, 50), randuniform(0, 1), randuniform(0, 1) d, e = randuniform(0.51, 1.0), randuniform(1, 50), f = randchoice([150,180,210,250,300]) 	215 216 217 218 219 220 221 222 223 224 225 226 227 228 229 230

We comparatively evaluated DODGE against AI tools which have been published in prior SE research. This paper does not compare DODGE against all other learners and all other hyperparameter optimizers (since such a comparison would not fit into a single paper). Instead, we use baselines as found in the SE literature about bad smell detection, predicting Github issue close time, bug report analysis, and defect prediction.

For example, for defect prediction, our classifiers come from a study by Ghotra et al. [25]. They found that the performance of dozens of data miners (applied to defect prediction) clustered into just few groups. By sampling a few algorithms from each group, we can explore the range of AI tools seen in defect prediction.

Clustering algorithms like EM [16] divide the data into related groups, then check the properties of each group. Another clustering method used in text mining, is Latent Dirichlet Allocation [11] that infers “topics” (commonly associated words). After documents are

scored according to how often they use some topic, a secondary classifier can then be used to distinguish the different topics.

Clustering algorithms like EM and LDA do not make use of any class variable. Naive Bayes classifiers [18], on the other hand, always divide the data on the class. New examples are then classified according to which class is most similar. Also, logistic regression fits the data to a particular parametric form (the logistic function).

Another learner that uses the class variable are decision tree algorithms [14, 65]. These learners divide the data using the attribute whose values most separate the classes. This learner then recurses on each division. Random Forests [13] build a “committee” of multiple decision trees, using different sub-samples of the data. Conclusions then come from some voting procedure across all the decision trees in the forest.

Standard clustering and decision tree algorithms base their analysis using the raw problem data. But what if some extra derived attribute is best at separating the classes? To address that issue, SVMs use some “kernel” to infer those extra dimension [12].

As for hyperparameter optimization, recall that these are tools to find a learner’s control settings (e.g., see Table 1). As mentioned above, we do not compare DODGE against all other hyperparameter optimization methods. Instead, we compare DODGE against prior hyperparameter optimization methods seen in SE literature. For example, for text mining of SE data, Panichella et al. [61] used genetic algorithm [27] (GA) to “evolve” a set of randomly generated control settings for SE text miners by repeating the following procedure, across many “generations”: (a) mutate a large population of alternate settings; (b) prune the worse performing settings; (c) combine pairs of the better, mutated options.

An alternative hyperparameter optimization strategy is the differential evolution [70] used by Wu et al. [23] and others [3]. DE generates mutants by interpolating between the better ranked settings. These better settings are kept in a “frontier list”. Differential evolution iterates over the frontier, checking each candidate against a new mutant. If the new mutant is better, it replaces the frontier item, thus improving the space of examples used for subsequent mutant interpolation.

Tantithamthavorn et al. [71] used a grid search for their hyperparameter optimization study. Grid search runs nested “for-loops” over the range of each control options. Fu et al. [24] found that for defect prediction, grid search ran 100 to 1000 times slower than DE.

3 DODGE: AN SE-BASED AI TOOL

This section describes DODGE, a tool described by Agrawal et al. [1] as an “SE-based” optimizer.

Given an ever evolving set of tools, languages, platforms, tasks, user expectations, development population, development practices, etc, we might expect that any prediction about an SE project will only ever be approximately accurate, i.e., within ϵ of the true value. Agrawal et al. reasoned that ϵ is not a problem to be solved, but a resource that could be exploited, as follows:

The RELAX heuristic: *Ignore anything less than ϵ .*

DODGE applies this RELAX heuristic to do hyperparameter optimization. To illustrate this process, consider the following example. We are exploring the hyperparameter space of Table 1 and we are

scoring each option via the recall and false alarm rate of the learner that uses those settings (for a definition of these scores, see §5).

Given performance goals with the range $0 \leq g \leq 1$, ϵ divides the performance output space into $(1/\epsilon)^g$ cells. For example, consider the $g = 2$ goals of recall and false alarm. These have minimum and maximum values of zero and one. Hence, if $\epsilon = 0.2$, then these scores divide into five regions (at 0.2, 0.4, 0.6, 0.8). As shown in Figure 1, these divided scores separate a two-dimensional plot of recall vs false alarm scores into $(1/0.2)^2 = 25$ cells. In those cells, green denotes good performance (high recall, low false alarm) and red denotes cells with relatively worse performance.




Figure 1: 25 outputs of size ϵ .

When billions of inputs (in Table 1) are mapped into the 25 cells of Figure 1, then many inputs are *redundant*, i.e., lead to the same outputs. The faster we can “dodge” redundant options, the faster we can move on to explore the other $(1/\epsilon)^g$ possible outputs.

To implement “dodging”, DODGE models Table 1 as a tree where all nodes have initial weights $w = 0$. Next, N_1 times, DODGE selects branches at random. We evaluate the options in a branch and if the resulting scores are within ϵ of any previous scores, then DODGE deprecates those options via $w = w - 1$. Else $w = w + 1$.

After that, DODGE freezes the selected branches found so far. N_2 times, DODGE then makes random selection to restrict any numeric ranges. When a range is initially evaluated, a random number $r = \text{random}(lo, hi)$ is selected and its weight $w(r)$ is set to zero. Subsequently, this weight is adjusted (as described above). When a new value is required (i.e., when the branch is evaluated again) then if the best, worst weights seen so far (in this range) are x, y (respectively) then we reset lo, hi to:

$$\text{IF } x \leq y \text{ THEN } lo, hi = x, (x + y)/2 \text{ ELSE } lo, hi = (x + y)/2, x$$

When ϵ is large, a few samples should suffice to find good results. Hence, Agrawal et al. [1] recommends $\epsilon = 0.2$ and $N_1 + N_2 = 30$.

DODGE can be recommended for two reasons. Firstly, for SE problems, DODGE’s optimizations are better than the prior state-of-the-art (evidence: see [1], and the rest of this paper). Secondly, DODGE achieves those results very quickly. Based on the default parameters suggested by Goldberg [27], Storn [70], and using some empirical results from Fu et al. [23], we can compute the number of times a hyperparameter optimizer would have to call a data miner. Assuming that, 25 times¹ we are analyzing 10 data sets, then hyperparameter optimization with grid search or genetic algorithms or differential evolution or DODGE would need to call a data miner 10^7 or 10^6 or 10^5 or 10^4 times (respectively).

4 BUT WHEN WILL DODGE FAIL?

This section describes the problem data used to evaluate DODGE.

Agrawal et al. say DODGE is an “SE-based” algorithm whose success depends on RELAX heuristic which, in turn, depends on the high variability of predictions for software. Their approach is so unusual that it is reasonable to ask “when won’t it work?”. To

¹Why 25? In a 5x5 cross-val experiment, the data set order is randomized five times. Each time, the data is divided into five bins. Then, for each bin, that bin becomes a test set of a model learned from the other bins.

amc	average method complexity
avg_cc	average McCabe
ca	afferent couplings
cam	cohesion amongst classes
cbm	coupling between methods
cbo	coupling between objects
ce	effluent couplings
dam	data access
dit	depth of inheritance tree
ic	inheritance coupling
lcom (lcom3)	2 measures of lack of cohesion in methods
loc	lines of code
max_cc	maximum McCabe
mfa	functional abstraction
moa	aggregation
noc	number of children
npm	number of public methods
rfc	response for a class
wmc	weighted methods per class
defects	Boolean: where defects found in bug-tracking

Table 2: Static code metrics for defect prediction. For details, see [42].

answer this question, the rest of this paper applies DODGE to 120 data sets, 83 of which come from SE while the others are from the AI literature. Agrawal et al. previously studied 16 of these data sets (and none of their data came from non-SE problems).

We will find that DODGE works very well for SE case studies and very badly for standard AI case studies. Later in this paper, we precisely characterize the kinds of data for which DODGE is not recommended.

All these pre-processed data sets are available on-line at <https://github.com/binded4review>.

4.1 Defect Prediction

Software developers are smart, but sometimes make mistakes. Hence, it is essential to test software before the deployment [8, 56, 59, 80]. Software bugs are not evenly distributed across the project [29, 39, 53, 60]. Hence, a useful way to perform software testing is to allocate most assessment budgets to the more defect-prone parts in software projects. Data miners can learn a predictor for defect proneness using, for e.g., the static code metrics of Table 2.

In a recent study, Rahman et al. [66] compared (a) static code analysis tools FindBugs, Jlint, and PMD and (b) static code defect predictors (which they called “statistical defect prediction”) built using logistic regression. They found no significant differences in the cost-effectiveness of these approaches. This result is interesting since code defect prediction can be quickly adapted to new languages by building lightweight parsers to extract static code metrics. The same is not true for static code analyzers that need much modification when languages update.

A recent survey of 395 practitioners from 33 countries and five continents [72] found that over 90% of the respondents were willing to adopt defect prediction techniques. Also, when Misirli et al. [53] built a defect prediction model for a telecommunications company, those models could predict 87% of files with defects. Those models also decreased inspection efforts by 72%, and reduced post-release defects by 44%.

Table 3 shows the static code data used in this paper. All these projects have multiple versions and we use older versions to predict the properties of the latest version. Note the fluctuating frequencies of the target class in the training and testing data (sometimes

Table 3: Defect prediction data from <http://tiny.cc/seacraft>. Uses metrics from Table 2.

Project	Training Data		Testing Data	
	Versions	% of Defects	Versions	% of Defects
Poi	1.5, 2.0, 2.5	426/936 = 46%	3.0	281/442 = 64%
Lucene	2.0, 2.2	235/442 = 53%	2.4	203/340 = 60%
Camel	1.0, 1.2, 1.4	374/1819 = 21%	1.6	188/965 = 19%
Log4j	1.0, 1.1	71/244 = 29%	1.2	189/205 = 92%
Xerces	1.2, 1.3	140/893 = 16%	1.4	437/588 = 74%
Velocity	1.4, 1.5	289/410 = 70%	1.6	78/229 = 34%
Xalan	2.4, 2.5, 2.6	908/2411 = 38%	2.7	898/909 = 99%
Ivy	1.1, 1.4	79/352 = 22%	2.0	40/352 = 11%
Synapse	1.0, 1.1	76/379 = 20%	1.2	86/256 = 34%
Jedit	3.2,4.0, 4.1,4.2	292/1257 = 23%	4.3	11/492 = 2%

Table 4: Issue tracking data (from <http://tiny.cc/seacraft>).

Dataset	No. of Documents	No. of Unique Words	Severe %
PitsA	965	155,165	39
PitsB	1650	104,052	40
PitsC	323	23,799	56
PitsD	182	15,517	92
PitsE	825	93,750	63
PitsF	744	28,620	64

increasing, sometimes decreasing), for e.g., xerces has target frequency changes between 16 to 74% while in jedit it changes from 23 to 2%. One of the challenges of doing data mining in such domains is finding learner settings that can cope with some wide fluctuations.

4.2 Text Mining Issue Reports

Many SE project artifacts come in the form of *unstructured text* such as word processing files, slide presentations, comments, Github issue reports, etc. According to White [73], 80% of business is conducted on unstructured data, 85% of all data stored is held in an unstructured format and unstructured data doubles every three months. Nadkarni and Yezhkova [57] say that 1600 Exabytes of data appears in unstructured sources and that each year, humans generate more unstructured artifacts than structured.

Mining such unstructured data is complicated by the presence of free form natural language which is semantically very complex and may not conform to any known grammar. In practice, text documents require tens of thousands of attributes (one for each word). For example Table 4 shows the number of unique words found in the issue tracking system for six NASA projects PitsA, PitsB, PitsC, etc. [44, 49]. Our PITS dataset contains tens to hundreds of thousands of words (even when reduced to unique words, there are still 10,000+ unique words). One other thing to note in Table 4 is that the target class frequencies are much higher than with defect prediction (median=60%).

For such large vocabulary problems, one of the tasks of a text miner is dimensionality reduction. See Table 1 for a list of dimensionality reduction pre-processing methods applied to the Table 4. After pre-processing, one of the learners from Table 1 was applied to predict for issue severity.

While these data mention five classes of severity, two of them comprise nearly all the examples. Hence, for this study we use the most common class and combine all the others into “other”. Agrawal

Table 5: Metrics used in Issue lifetime data

Commit	Comment	Issue
nCommitsByActorsT	meanCommentSizeT	issueCleanedBodyLen
nCommitsByCreator	nComments	nIssuesByCreator
nCommitsByUniqueActorsT		nIssuesByCreatorClosed
nCommitsInProject		nIssuesCreatedInProject
nCommitsProjectT		nIssuesCreatedInProjectClosed
		nIssuesCreatedProjectClosedT
		nIssuesCreatedProjectT
Misc.	nActors, nLabels, nSubscribedByT	

et al. [1] showed that using Table 1, they could auto-configure classifiers to better predict for this binary severity problem.

4.3 Issue Lifetime Estimation

Issue tracking systems collect information about system failures, feature requests, and system improvements. Based on this information and actual project planing, developers select the issues to be fixed.

Predicting the time it may take to close an issue has multiple benefits for the developers, managers, and stakeholders involved in a software project. Such predictions helps software developers to better prioritize work. For an issue close time prediction generated at issue creation time can be used, for example, to auto-categorize the issue or send a notification if it is predicted to be an easy fix. Also, such predictions helps managers to effectively allocate resources and improve consistency of release cycles. Lastly, such predictions helps project stakeholders understand changes in project timelines.

Such predictions can be generated via data mining. Rees-Jones et al. [68] analyzed the Giger et al. [26] data using Hall's CFS feature selector [28] and the C4.5 decision tree learner [65]. They found that the attributes of Table 5 could be used to generate very accurate predictions for issue lifetime. Table 6 shows information about the nine projects used in the Rees-Jones study. Note here that the target class frequencies vary greatly from 2 to 42%.

4.4 Bad Code Smell Detection

According to Fowler [21], bad smells (i.e., code smells) are "a surface indication that usually corresponds to a deeper problem". Studies suggest a relationship between code smells and poor maintainability or defect proneness [76, 77, 81]. Research on software refactoring endorses the use of code-smells as a guide for improving the quality of code as a preventative maintenance. Consequently, code smells are captured by popular static analysis tools, like PMD, Check-Style, FindBugs, and SonarQube. Much recent progress has been made towards adopting data mining AI tools to classify code smells. Kreimer [40] proposes an adaptive detection to combine known methods for finding design flaws. Khomh et al. [36] proposed a Bayesian approach to detect occurrences of the Blob antipattern on open-source programs. Khomh et al. [37] also proposed an approach to build Bayesian Belief Networks. Yang et al. [78] studied the judgment of individual users by applying machine learning algorithms on code clones.

Recently, Fontana et al. [5] considered 74 systems in data mining analysis. Table 7 shows the data used in that analysis. This corpus comes from 11 systems written in Java, characterized by different sizes and belonging to different application domains. The authors computed a large set of object-oriented metrics belonging to a class,

Table 6: Issue Lifetime Estimation Data from [68]

Project Name	Dataset	# of instances		# metrics (see Table 5).
		Total	Closed (%)	
camel	1 day	5056	698 (14.0)	18
	7 days		437 (9.0)	
	14 days		148 (3.0)	
	30 days		167 (3.0)	
	90 days		298 (6.0)	
	180 days		657 (13.0)	
cloudstack	1 day	1551	2052 (41.0)	18
	7 days		658 (42.0)	
	14 days		457 (29.0)	
	30 days		101 (7.0)	
	90 days		107 (7.0)	
	180 days		133 (9.0)	
cocoan	1 day	2045	65 (4.0)	18
	7 days		92 (4.0)	
	14 days		32 (2.0)	
	30 days		45 (2.0)	
	90 days		86 (4.0)	
	180 days		51 (3.0)	
node	1 day	6207	73 (3.5)	18
	7 days		2426 (39.0)	
	14 days		1800 (29.0)	
	30 days		521 (8.0)	
	90 days		453 (7.0)	
	180 days		552 (9.0)	
deeplearning	1 day	1434	180 (3.0)	18
	7 days		931 (65.0)	
	14 days		214 (15.0)	
	30 days		76 (5.0)	
	90 days		72 (5.0)	
	180 days		69 (5.0)	
hadoop	1 day	12191	39 (3.0)	18
	7 days		40 (0.0)	
	14 days		65 (1.0)	
	30 days		107 (1.0)	
	90 days		396 (3.0)	
	180 days		1743 (14.0)	
hive	1 day	5648	2133 (17.5)	18
	7 days		18 (0.0)	
	14 days		22 (0.0)	
	30 days		58 (1.0)	
	90 days		178 (3.0)	
	180 days		1050 (19.0)	
ofbiz	1 day	6177	1356 (24.0)	18
	7 days		1440 (25.0)	
	14 days		1515 (25.0)	
	30 days		1169 (19.0)	
	90 days		467 (8.0)	
	180 days		477 (8.0)	
qqid	1 day	5475	574 (9.0)	18
	7 days		469 (7.5)	
	14 days		402 (6.5)	
	30 days		203 (4.0)	
	90 days		188 (3.0)	
	180 days		84 (2.0)	
365 days	1 day	5475	178 (3.0)	18
	7 days		558 (10.0)	
	14 days		860 (16.0)	
	30 days		531 (10.0)	
	90 days			
	180 days			

Table 7: Bad code smell detection data from [5]

Nature	Dataset	No. of instances	No. of attributes	Smelly %
Method	Feature Env	109	82	45
Method	Long Method	109	82	43.1
Class	God Class	139	61	43.9
Class	Data Class	119	61	42

method, package, and project level. A detailed list of metrics are available in appendices of [5]. Note in Table 7, how the target class frequencies are all around 43%,

Table 8: NON-SE problems: 37 UCI Datasets statistics.

Area	Dataset	# of instances	# of attributes	Class %
Computer	optdigits	1143	64	50
Physical	satellite	2159	36	28
Physical	climate-sim	540	18	91
Financial	credit-approval	653	15	45
Medicine	cancer	569	30	37
Business	shop-intention	12330	17	15
Computer Vision	image	660	19	50
Life	covtype	12240	54	22
Computer	hand	29876	15	47
Social	drug-consumption	1885	30	23
Environment	biodegrade	1055	41	34
Social	adult	45222	14	25
Physical	crowdsourcing	1887	28	24
Medicine	blood-transfusion	748	4	24
Financial	credit-default	30000	23	22
Medicine	cervical-cancer	668	33	7
Social	autism	609	19	30
Marketing	bank	3090	20	12
Financial	bankrupt	4769	64	3
Financial	audit	775	25	39
Life	contraceptive	1473	9	56
Life	mushroom	5644	22	38
Computer	pendigits	2288	16	50
Security	phishing	11055	30	56
Automobile	car	1728	6	30
Medicine	diabetic	1151	19	53
Physical	hepmass	2000	27	50
Physical	htru2	17898	8	9
Computer	kddcup	3203	41	69
Automobile	sensorless-drive	10638	48	50
Physical	waveform	3304	21	50
Physical	annealing	716	10	13
Medicine	cardiotocography	2126	40	22
Physical	shuttle	54489	9	16
Electrical	electric-stable	10000	12	36
Physical	gamma	19020	10	35
Medicine	liver	579	10	72

4.5 Non-SE Problems

The UCI machine learning repository [7, 17, 22] was created in 1987 to foster experimental research in machine learning. To say the least, this repository is commonly used by industrial and academic researchers (evidence: the 2007, 2010, and 2017 version of the repository are cited 4020, 3179 and 4179 times respectively [7, 17, 22]). Many of AI tools used in SE were certified using data from UCI. This repository holds 100s of data mining problems from many problem areas including engineering, molecular biology, medicine, finance and politics. Using a recent state-of-the-art machine learning paper [74] we identified 37 UCI data sets that machine learning researchers often use in their analysis (see Table 8).

One issue with comparing Table 8 to the SE problems is that the former often have $N > 2$ classes whereas the SE problems use binary classification. Also, sometimes, the SE data exhibits large class imbalances (where the target is less than 25% of the total). Such imbalances are acute in the issue lifetime data in Table 6 but it also appears sometimes in the test data of Table 3.

We considered various ways to remove this threat to validity including (a) clustering and sub-sampling each cluster; (b) some biased sampling of the UCI data. In then end, we adopted a very simple method (lest anything more complex introduced its own biases). For each UCI dataset, we selected:

- The UCI rows from the most frequent and rarest class;
- And declared that the UCI rarest class is the target class.

5 EXPERIMENTS AND RESULTS

To comparatively evaluate the performance of AI tools, we need performance measures (§5.1), appropriate experimental rigs (§5.2), and statistical tests (§5.3).

5.1 Performance Measures

$D2h$, or “distance to heaven”, shows how close a classifier falls to “heaven” (where recall=1 and false alarms (FPR)=0) [15]:

$$\text{Recall} = \text{TruePositives}/(\text{TruePositives} + \text{FalseNegatives}) \quad (1)$$

$$\text{FPR} = \text{FalsePositives}/(\text{FalsePositives} + \text{TrueNegatives}) \quad (2)$$

$$d2h = \left(\sqrt{(1 - \text{Recall})^2 + (0 - \text{FPR})^2} \right) / \sqrt{2} \quad (3)$$

Here, the $\sqrt{2}$ term normalizes $d2h$ to the range zero to one.

For defect prediction, $Popt(20)$ comments on the inspection effort required *after* a defect predictor is triggered. $Popt(20) = 1 - \Delta_{opt}$, where Δ_{opt} is the area between the effort (code-churn-based) cumulative lift charts of the optimal learner and the proposed learner. To calculate $Popt(20)$, we divide all the code modules into those predicted to be defective (D) or not (N). Both sets are then sorted in ascending order of lines of code. The two sorted sets are then laid out across the x-axis, with D before N . On such a chart, the y-axis shows what percent of the defects would be recalled if we traverse the code sorted that x-axis order. Following from Ostrand et al. [60], $Popt$ is reported at the 20% point. Further, following Kamei, Yang et al. [35, 55, 79] we normalize $Popt$ using: $Popt(m) = 1 - \frac{S(optimal) - S(m)}{S(optimal) - S(worst)}$ where $S(optimal)$, $S(m)$ and $S(worst)$ represent the area of curve under the optimal learner, proposed learner, and worst learner. Note that the worst model is built by sorting all the changes according to the actual defect density in ascending order. After normalization, $Popt(20)$ (like $d2h$) has the range zero to one. Note that *larger* values of $Popt(20)$ are *better*, but *smaller* values of $d2h$ are *better*.

5.2 Experimental Rigs

Jimenez et al. [33] recommended that train/test data be labelled in their natural temporal sequence. E.g. for defect data, we apply training and hyperparameter optimization to the prior versions, then tested on the latest version. We call this approach **RIG0**.

When such temporal markers are absent, we applied a cross-val method (which is also standard in literature, for e.g., see [79]). Given one data set and N possible treatments, then 25 times we use 80% of the data (selected at random) for training and hyperparameter optimization, then use the remaining 20% of the data for testing. We call this approach **RIG1**.

5.3 Statistical Tests

When comparing results from two samples, we need a statistical significance test (to certify that the distributions are indeed different) and an effect size test (to check that the differences are more than a “small effect”). Here, we used tests which have been past peer reviewed in the literature [2, 3]. Specifically, we use Efron’s 95% confidence bootstrap procedure [19] and the A12 effect test endorsed by Acuri & Briand in their ICSE paper [6].

5.4 Results

In the following, when we say “DODGE”, that is shorthand for DODGE using Table 1 with $N_1 + N_2 = 15$, $\epsilon = 0.2$. Also, when we say “DODGE performed better”, we mean that, according to a 95% bootstrap and the A12 test, DODGE performed significantly better by more than a small effect.

5.4.1 Defect Prediction. Table 9 shows which AI tools found best predictors for defects, using the data of §4.1.

One issue with defect prediction is the occasional presence of class imbalance. If the target class is not common (as in camel, ivy, jedi and to a lesser extent velocity and synapse), it can be difficult for a data mining algorithm to generate a model that can locate it. Researchers have used class balancing techniques such as SMOTE to address this problem [3].

Table 9 study compares DODGE versus methods selected from prior state-of-the-art SE papers. An ICSE'18 paper [3] reported that hyperparameter tuning (using DE) of SMOTE usually produces best results (and this result holds across multiple learners, applied after class rebalancing). We used SMOTE tuning (for data-processing) plus learners taken from Ghotra et al. [25] (who found that the performance of dozens of data miners can be clustered into just a few groups). We used learners sampled across those clusters (Random Forests, CART, SVM, KNN ($k = 5$), Naive Bayes, Logistic Regression).

Table 9 results were generated using **RIG0** with *d2h* and *Popt(20)* as the performance goal. DODGE performed statistically better than the prior state-of-the-art in sixteen out of twenty data sets.

5.4.2 Text Mining Issues Reports. Table 10 shows which AI tools found best predictors for the data of §4.2. In this study, all data were preprocessed using the usual text mining filters [20]. We implemented stop words removal using NLTK toolkit [10] (to ignore very common short words such as “and” or “the”). Next, Porter’s stemming filter [64] was used to delete uninformative word endings (e.g., after performing stemming, all the following words would be rewritten to “connect”: “connection”, “connections”, “connective”, “connected”, “connecting”).

Table 10 study compares DODGE versus methods seen in prior state-of-the-art SE papers: specifically, SVM plus Latent Dirichlet allocation [11] with hyperparameter optimization via differential evolution [2] or genetic algorithms [61].

Table 9: Ten defect prediction results. Smote+ means SMOTE+ DE tuning + best of the Ghotra’15 learners.

data	Best tool	
	D2h	Popt
Poi	Dodge	Smote+
Lucene	Dodge	Dodge
Camel	Dodge	Dodge
Log4j	Smote+	Dodge
Xerces	Dodge	Dodge
Velocity	Dodge	Dodge
Xalan	Smote+	Dodge
Ivy	Dodge	Dodge
Synapse	Dodge	Dodge
Jedit	Smote+	Dodge

Table 10: Six text mining results.

Data	Best tool
PitsA	Dodge
PitsB	LDA + DE
PitsC	Dodge
PitsD	Dodge
PitsE	Dodge
PitsF	Dodge

Table 11: Sixty three issue lifetime prediction results. DODGE loses to random forests in colored cells.

Data	Days till closed						
	> 365	< 180	< 90	< 30	< 14	< 7	< 1
cloudstack	Dodge	Dodge	Dodge	RF	RF	Dodge	RF
node	Dodge	Dodge	Dodge	Dodge	RF	Dodge	Dodge
deeplearning	Dodge	Dodge	Dodge	Dodge	Dodge	Dodge	RF
cocoon	RF	RF	Dodge	Dodge	Dodge	Dodge	Dodge
ofbiz	Dodge	Dodge	Dodge	Dodge	Dodge	Dodge	Dodge
camel	RF	Dodge	Dodge	RF	RF	Dodge	Dodge
hadoop	Dodge	Dodge	Dodge	RF	RF	Dodge	RF
qpud	Dodge	Dodge	Dodge	Dodge	Dodge	Dodge	Dodge
hive	RF	RF	Dodge	Dodge	Dodge	Dodge	Dodge
DODGE wins	6/9	7/9	9/9	6/9	5/9	9/9	6/9

Table 10 results were generated using **RIG1** with *d2h* as the performance goal. In these results, DODGE performed statistically better than the prior state-of-the-art (in five out of six data sets).

5.4.3 Issue Lifetime Estimation. Table 11 shows which AI tools found best predictors, using the data of §4.3. The table compares DODGE versus the methods in a recent study on issue lifetime estimation [68], i.e., feature selection with the Correlation Feature Selection algorithm [28] followed by classification using Random Forests.

Table 11 results were generated using **RIG1** with *d2h* as the performance goal. In these results, DODGE performed statistically better than prior work (in 47/63=75% of the datasets).

5.4.4 Bad Code Smell Detection. Figure 2 shows which AI tools found best predictors, using the data of §4.4. The figure compares DODGE versus bad smell detectors from a TSE'18 paper [41] that studied bad smells. The TSE article used Decision Trees (CART), Random Forests, Logistic Regression and KNN($k = 5$). To the best




Figure 2: Four bad smell prediction results. Boxes indicate 75th-25th range seen in 25 repeats. Whiskers extend from min to max. Horizontal line in the middle of each box indicates median value. In this figure, lower values are better.

Table 12: Thirty seven results from non-SE problems. Colored cells show where DODGE performs best.

Data	Best tool	Data	Best tool
optdigits	RF	satellite	RF
climate-sim	SVM	credit-approval	Dodge
cancer	SVM	shop-intention	RF
image	RF	covtype	RF
hand	RF	drug-consumption	Dodge
biodegrade	RF	adult	RF
crowdsource	RF	blood-transfusion	Dodge
credit-default	SVM	cervical-cancer	Dodge
autism	RF	bank	SVM
bankrupt	Dodge	audit	RF
contraceptive	SVM	mushroom	RF
pendigits	RF	phishing	RF
car	RF	diabetic	SVM
hepmass	RF	htru2	SVM
kddcup	RF	sensorless-drive	RF
waveform	SVM	annealing	RF
cardiotocography	RF	shuttle	RF
electric-stable	RF	gamma	RF
liver	Dodge		

of our knowledge, there has not been any prior case study applying hyperparameter optimizer to bad smell prediction.

Figure 2 results were generated using **RIG1** with *d2h* as the performance goal. In those results, DODGE has the same median performance as prior work for two data sets (FeatureEnvy and GodClass) and performed statistically better than the prior state-of-the-art (for DataClass and LongMethod). That is, compared to the other algorithms used in this study, DODGE statistically performs as well or better than anything else.

5.4.5 Non SE-Problems. The results so far support DODGE’s approach to hyperparameter optimization. But all the above problems come from the SE domain. When we turn to non-SE problems, a very different pattern was observed.

Table 12 shows which AI tools found best predictors for the 37 different problems from §4.5.

In Table 12, DODGE was compared against standard data miners (CART, Random Forests, Logistic Regression and KNN($k = 5$)). Table 12 results were generated using **RIG1** with *d2h* as the performance goal. Each cell of that table lists the best performing learner. Note that despite its use of hyperparameter optimization (which should have given some advantage) DODGE performs very badly (only succeeds in 6/31 problems).

6 META-ANALYSIS

This section addresses the puzzle raised by the above results. DODGE works very well for SE problems and works very poorly for non-SE problems. Why?

6.1 Methodology

This section describes our method for learning what kind of problems works best for SE/AI-based tools.

- (1) Given M treatments for data, divide those treatments into those from SE and those from elsewhere. In our case, DODGE is the SE-based tool. Also, as discussed in §2.2, grid search, DE, GA, SVM, Random Forests, decision tree learners, logistic regression, Naive Bayes, LDA, and EM are representative of standard AI.
- (2) Apply those M treatments to P data sets (some SE, some from elsewhere). In the case of this paper, our data was described above.
- (3) (Optional): As done in this paper, we recommend guiding that exploration by adopting the methods seen in recent publications in high-profile SE venues (since such methods have some pedigree in the literature).
- (4) Using statistical methods and experimental rigs appropriate to those problems, build a table of data showing what AI tools work best for each data sets. In that table:
 - For the dependent class, use the name of the best AI tool.
 - For the independent variables, use some data characteristics shared by all your problems. For this study, we used the independent attributes described in the next section.
- (5) Run a decision tree learner (CART or C4.5) over that table.
- (6) Read the generated tree to learn what kind of data is best processed by general AI tools or SE-based tools.

6.2 Independent Variables for the Meta-Analysis

To understand when DODGE performs well/fail, we need some characterization of the data sets that holds across most of the above data. All of the above problems are stored in tables of data and can be described via:

- N = Number of samples (rows of example instances);
- F = Frequency of our selected target class;

Of course, our data has more independent information than these two points. For example, we tried using some other attributes (the percentile distributions of the expected value of the attribute entropy). But when we applied the CFS [28] feature selector, all those other independent attributes kept being deleted. Hence, to build a rule for determining when to use/avoid DODGE, we used (N , F).

The data generated in this way is shown at http://tiny.cc/variance_defense.

6.3 Meta-Analysis Results

For this analysis, anything not labelled as “DODGE” was renamed “general AI tool”. The following rule was generated. The number of brackets after each conclusion shows the percentage of the data that falls to each branch.

```

IF F <= 21.15%
THEN
  IF N <= 12215
  THEN use DODGE (55)
  ELSE use general AI tools (3)
THEN
  use general AI tools (42)

```

In summary, if the target class is not rare (more than 21%) then use general AI tools. Otherwise, except for one infrequent case (3% of our data), use DODGE.

Since this rule was learned via data mining, it is not a categorical statement that, always, anything confirming to one of its branches

will 100% benefit best from DODGE or some general AI tool. There are certainly counter-examples to this rule in our data (particularly with the text mining and bad smell data where the target classes are quite common). Hence, in a 5-way cross-validation experiment, this rule did not report 100% recalls or precisions for selecting which tools:

class	recall	precision	false alarm
DODGE	77	72	28
General AI	72	70	27

Nevertheless, given the complexity of the phenomena being explored here, the observed recall, precision, and false alarm rates are remarkably good. Also, as discussed in the next section, we found this rule insightful.

6.4 Discussion

We posted the rule learned above to a machine learning discussion forum and received back the following insight. For many decades, the AI community has been using the UCI repository as a source of data for certifying their algorithms. Several of the data sets of Table 8 have appeared in publications dating back at least to 1995 i.e., nearly 25 years ago [52]. Quirks of that data have hence led to the development of algorithms that are skewed toward particular kinds of data.

To say that another way, the algorithms that we have called “general AI tools” may not be “general” at all. Rather, they may be tools that are powerful in their home domain but need to be used with care if applied to new domains such as software engineering. Many data mining problems in software engineering target infrequent phenomena. Hence, when the target class is relatively uncommon (say, less than 20%), we can do better than just using AI tools built and tested for data with much more frequent classes.

As to the connection of infrequent classes to the RELAX heuristic of §3, it seems reasonable to assume that the more common the target class, then:

- The easier it is to learn about the target;
- The finer the control the learner has about its generated models;
- The smaller the ϵ inaccuracy measure in the predictions;
- The fewer the redundant tuning options;
- And the less likely that the limited sampling of DODGE will find good tunings;

On the other hand, when the target class becomes less frequent, then:

- It is harder to find the target;
- The larger the observed ϵ in the results;
- The greater the number of redundant tunings;
- The more likely that DODE will work.

That is, since software engineering often deals with relatively infrequent target classes, then we should expect to see a large ϵ uncertainty in our conclusions.

7 THREATS TO VALIDITY

7.1 Sampling Bias

Sampling bias threatens any classification experiment since what matters for some data sets may or may not hold for others.

But in our case, sampling bias may be mitigated since we applied our frameworks to many data sets. In fact, our reading of the literature is that the above study uses much more data than most other publications. Also, we assert we did not “cherry pick” our data sets. All the non-SE data from [74] were applied here. As to the SE data, we used everything we could access in the time frame of this production of this paper.

That said, in future work, it would be important to check the results of this paper against yet more data sets from yet more problems from SE and elsewhere.

7.2 Learner Bias

When comparing DODGE against other AI tools, we did not explore all other AI tools. As stated above, such a comparison would not fit into a single paper. Instead, we use baselines taken from SE literature about bad smell detection, predicting Github issue close time, bug report analysis, and defect prediction. From that work we used tools that have some pedigree in the literature [3, 25, 61, 71].

7.3 Evaluation Bias

This paper used two performance measures, i.e., P_{opt} and $dist2heaven$ and many others exist [34, 46, 50]. Note that just because other papers use a particular evaluation bias, then it need not follow that it must be applied in all other papers. For example, precision is a widely used evaluation method even though it is known to perform badly when data sets have rare target classes [45].

7.4 Order Bias

For the performance evaluation part, the order that the data trained and predicted can affects the results. To mitigate for order bias, we used a cross-validation procedure that (multiple times) randomizes the order of the data.

7.5 Construct Validity

At various stages of data collection by different researchers, they must have made engineering decisions about what attributes to be extracted from Github for Issue lifetime data sets, or what object-oriented metrics need to be extracted. While this can inject issues of construct validity, we note that the data sets used here have also appeared in other SE publications, i.e., the class labels used here have verified by other researchers. So at the very least, using this data, we have no more construct validity bias than other researchers.

7.6 Statistical Validity

To increase the validity of our results, we applied two statistical tests, bootstrap and the a12 effect size test. Both of these are non-parametric tests so the above results are not susceptible to issues of parametric bias.

7.7 External Validity

One threat to external validity is that this article compares DODGE against existing baselines for traditional machine learning algorithms. We do not compare our new approach against the kinds of optimizers we might find in the search-based SE literature [63]. The reason for this is that the main point of this paper is to document a previously unobserved feature of the output space of software

analytics. It is an open question whether or not DODGE is the best way to explore output space. In order to motivate the community to explore that space, some work must demonstrate its existence and offer baseline results that, using the knowledge of output space, it is possible to do better than past work. We hope that this paper provides that motivation.

8 FUTURE WORK

Further to the discussion of in the *Threats to Validity* section, the analysis of this paper should be extended with studies on different data sets and different AI tools.

Another useful extension to the above would be to explore problems with three or more goals (e.g., reduce false alarms while at the same time improving precision and recall). Currently, we have only explored classification tasks. DODGE needs to be applied for other tasks such as regression task.

Right now, DODGE only deprecate tunings that lead to similar results. Another approach would be to also depreciate tunings that lead to similar *and worse* results (perhaps to rule out larger parts of the output space, sooner).

Further, for pragmatically reasons it would be useful if the Table 1 list could be reduced to a smaller, faster to run, set of learners. That is, here we would select learners that run fastest while generating the most variable kinds of models.

Moving beyond DODGE, it would be useful to test the generality of the meta-analysis of §6. For example, Helendoorn and Devanbu [31] report cases where a standard AI tool (deep learning) is defeated by their own SE-based method (based on software “naturalness”). To repeat the analysis of this paper, it would be useful to collect a large library of problems where deep learning does/does not work best, then look for problem data features that predict for the success of deep learning.

9 CONCLUSION

One way to view SE is that it is just a problem domain where general AI tools can be deployed. This paper takes another view. We say software engineering is different and that those differences mean that best results arise from tools specialized to the particulars of software engineering.

For example, SE problem data is often about infrequent issues, for e.g., the occasional defect, the rarely exploited security violation, the requirement that only holds for one small part of a system. But as shown by the rule in §6.3, standard AI tools work best when the target is relatively more frequent. That is:

- SE problem data may be inherently different to the problem data used to develop standard AI tools.
- And that difference can be used to (e.g.,) find better ways to build predictions for SE classification problems such as defect prediction, issue text mining, Github issue lifetime estimation, and bad code smell detection.

The tool used to find these better predictions, called DODGE, runs orders of magnitude faster than other methods (see calculations at the end of §5). We conjecture that DODGE works better and faster than those other tools since those other general AI tools do not appreciate the simplicity of the output space of SE problems (Figure 1):

- AI tools such as grid search, differential evolution, or genetic algorithms waste much CPU as they struggle to cover billions of tuning options like Table 1 (most of which yield indistinguishably different results).
- This is a waste of effort since when billions of possibilities (Table 1) are mapped into just a few outputs (Figure 1), then many of those possibilities are *redundant*; i.e., they lead to the same outputs.

DODGE knows that the faster we move away from these redundant possibilities, the faster we can move on to explore the remaining options. Hence, one lesson from the DODGE research is to say:

- Stop treating conclusion uncertainty as a problem;
- Instead, treat large conclusion uncertainty as a resource that can be used to simplify the analysis of software systems.

In conclusion, it is not good enough to just take AI tools developed elsewhere, then apply them without modification to SE problems. Understanding SE is a different problem to understanding other problems that are more precisely controlled and restrained. The algorithms that might be called “general AI tools” may not be “general” at all. Rather, they are tools which are powerful in their home domain but need to be used with care if applied to some new domains like software engineering. Software engineers need to develop AI tools that are better suited to the particulars of SE problems. As shown in this paper:

- Such new algorithms can exploit the peculiarities of SE data to dramatically simplify the analysis of software systems.
- Further, as shown in §6.3, it is possible to define a rule that precisely defines when *not* to use those SE-based tools.

ACKNOWLEDGEMENTS

This work was partially funded by blinded for review.

REFERENCES

- [1] Amritanshu Agrawal, Wei Fu, Di Chen, Xipeng Shen, and Tim Menzies. 2019. How to “DODGE” Complex Software Analytics? *CoRR* abs/1902.01838 (2019). arXiv:1902.01838 <http://arxiv.org/abs/1902.01838>
- [2] Amritanshu Agrawal, Wei Fu, and Tim Menzies. 2018. What is wrong with topic modeling? And how to fix it using search-based software engineering. *Information and Software Technology* (2018).
- [3] Amritanshu Agrawal and Tim Menzies. 2018. Is better data better than better data miners?: on the benefits of tuning SMOTE for defect prediction. In *International Conference on Software Engineering*.
- [4] Amritanshu Agrawal, Tim Menzies, Leandro L. Minku, Markus Wagner, and Zhe Yu. 2018. Better Software Analytics via “DUO”: Data Mining Algorithms Using/Used-by Optimizers. *CoRR* abs/1812.01550 (2018). <http://arxiv.org/abs/1812.01550> Available online at <https://arxiv.org/abs/1812.01550>.
- [5] Francesca Arcelli Fontana, Mika V. Mäntylä, Marco Zanoni, and Alessandro Marino. 2016. Comparing and experimenting machine learning techniques for code smell detection. *Empirical Software Engineering* (2016).
- [6] Andrea Arcuri and Lionel Briand. 2011. A Practical Guide for Using Statistical Tests to Assess Randomized Algorithms in Software Engineering. In *International Conference on Software Engineering*.
- [7] Arthur Asuncion and David Newman. 2007. UCI machine learning repository.
- [8] Earl T Barr et al. 2015. The oracle problem in software testing: A survey. *IEEE transactions on software engineering* (2015).
- [9] Dave Binkley, Dawn Lawrie, and Christopher Morrell. 2018. The Need for Software Specific Natural Language Techniques. *Empirical Softw. Engg.* 23, 4 (Aug. 2018), 2398–2425. <https://doi.org/10.1007/s10664-017-9566-5>
- [10] Steven Bird. 2006. NLTK: the natural language toolkit. In *Proceedings of the COLING/ACL on Interactive presentation sessions*.
- [11] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022. <http://dl.acm.org/citation.cfm?id=944919.944937>

- 1359 [12] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. 1992. A Training
1360 Algorithm for Optimal Margin Classifiers. In *Proceedings of the Fifth Annual*
1361 *Workshop on Computational Learning Theory (COLT '92)*. ACM, New York, NY,
USA, 144–152. <https://doi.org/10.1145/130385.130401>
- 1362 [13] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- 1363 [14] Leo Breiman. 2017. *Classification and regression trees*. Routledge.
- 1364 [15] Di Chen et al. 2018. Applications of Psychological Science for Actionable Analyt-
1365 ics. *Foundations of Software Engineering* (2018).
- 1366 [16] A. P. Dempster, N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from in-
1367 complete data via the EM algorithm. *JOURNAL OF THE ROYAL STATISTICAL*
1368 *SOCIETY, SERIES B* 39, 1 (1977), 1–38.
- 1369 [17] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- 1370 [18] Richard O. Duda, Peter E. Hart, and David G. Stork. 2000. *Pattern Classification*
1371 *(2Nd Edition)*. Wiley-Interscience, New York, NY, USA.
- 1372 [19] B. Efron and R. J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Chapman &
1373 Hall, New York.
- 1374 [20] R-SANGER Feldman. 2006. *J. The Text Mining Handbook*. New York: Cambridge
1375 University Press.
- 1376 [21] M. Fowler, K. Beck, J. Brant, W. Opdyke, and D Roberts. 1999. *Refactoring:*
1377 *Improving the Design of Existing Code*. Addison-Wesley Longman.
- 1378 [22] Andrew Frank. 2010. UCI machine learning repository. <http://archive.ics.uci.edu/ml> (2010).
- 1379 [23] Wei Fu, Tim Menzies, and Xipeng Shen. 2016. Tuning for software analytics: Is it
1380 really necessary? *Information and Software Technology* 76 (2016), 135–146.
- 1381 [24] Wei Fu, Vivek Nair, and Tim Menzies. 2016. Why is Differential Evolution Better
1382 than Grid Search for Tuning Defect Predictors? *CoRR* abs/1609.02613 (2016).
1383 arXiv:1609.02613 <http://arxiv.org/abs/1609.02613>
- 1384 [25] Baljinder Ghotra, Shane McIntosh, and Ahmed E Hassan. 2015. Revisiting the im-
1385 pact of classification techniques on the performance of defect prediction models.
1386 In *International Conference on Software Engineering*.
- 1387 [26] Emanuel Giger, Martin Pinzger, and Harald Gall. 2010. Predicting the fix time
1388 of bugs. In *Proceedings of the 2nd International Workshop on Recommendation*
1389 *Systems for Software Engineering*. ACM, 52–56.
- 1390 [27] David E Goldberg and John H Holland. 1988. Genetic algorithms and machine
1391 learning. *Machine learning* 3, 2 (1988), 95–99.
- 1392 [28] Mark A Hall and Geoffrey Holmes. 2002. Benchmarking attribute selection
1393 techniques for discrete class data mining. (2002).
- 1394 [29] Maggie Hamill and Katerina Goseva-Popstojanova. 2009. Common trends in
1395 software fault and failure data. *IEEE Transactions on Software Engineering* (2009).
- 1396 [30] Ahmed E Hassan. 2008. The road ahead for mining software repositories. In *2008*
1397 *Frontiers of Software Maintenance*. IEEE, 48–57.
- 1398 [31] Vincent J Hellendoorn and Premkumar Devanbu. 2017. Are deep neural networks
1399 the best choice for modeling source code?. In *Foundations of Software Engineering*.
1400 ACM.
- 1401 [32] Abram Hindle, Earl T Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu.
1402 2012. On the naturalness of software. In *2012 34th International Conference on*
1403 *Software Engineering (ICSE)*. IEEE, 837–847.
- 1404 [33] Matthieu Jimenez, Renaud Rwemalika, Mike Papadakis, Federica Sarro, Yves
1405 Le Traon, and Mark Harman. 2019. The Importance of Accounting for Real-
1406 world Labelling when Predicting Software Vulnerabilities. In *Proceedings of the*
1407 *2019 27th ACM Joint Meeting on European Software Engineering Conference and*
1408 *Symposium on the Foundations of Software Engineering (ESEC/FSE 2019)*. ACM,
1409 New York, NY, USA, 695–705. <https://doi.org/10.1145/3338906.3338941>
- 1410 [34] Magne Jorgensen. 2004. Realism in assessment of effort estimation uncertainty:
1411 It matters how you ask. *IEEE Transactions on Software Engineering* (2004).
- 1412 [35] Yasutaka Kamei, Emad Shihab, Bram Adams, Ahmed E Hassan, Audris Mockus,
1413 Anand Sinha, and Naoyasu Ubayashi. 2013. A large-scale empirical study of
1414 just-in-time quality assurance. *IEEE Transactions on Software Engineering* (2013).
- 1415 [36] F. Khomh, S. Vaucher, Y. G. Gueheneuc, and H. Sahraoui. 2009. A Bayesian
1416 Approach for the Detection of Code and Design Smells. In *International Conference*
1417 *on Quality Software*.
- 1418 [37] Foutse Khomh, Stephane Vaucher, Yann-Gael Gueheneuc, and Houari Sahraoui.
1419 2011. BDTEX: A GQM-based Bayesian approach for the detection of antipatterns.
1420 *Journal of Systems and Software* (2011).
- 1421 [38] Miryung Kim, Thomas Zimmermann, Robert DeLine, and Andrew Begel. 2016.
1422 The emerging role of data scientists on software development teams. In *Proceed-*
1423 *ings of the 38th International Conference on Software Engineering*. ACM, 96–107.
- 1424 [39] A Güneş Koru, Dongsong Zhang, Khaled El Emam, and Hongfang Liu. 2009. An
1425 investigation into the functional form of the size-defect relationship for software
1426 modules. *IEEE Transactions on Software Engineering* (2009).
- 1427 [40] Jochen Kreimer. 2005. Adaptive detection of design flaws. *Electronic Notes in*
1428 *Theoretical Computer Science* (2005).
- 1429 [41] Rahul Krishna and Tim Menzies. 2018. Bellwethers: A Baseline Method For
1430 Transfer Learning. *IEEE Transactions on Software Engineering* (2018).
- 1431 [42] Rahul Krishna, Tim Menzies, and Lucas Layman. 2017. Less is more: Minimizing
1432 code reorganization using XTREE. *Information and Software Technology* (mar
1433 2017). <https://doi.org/10.1016/j.infsof.2017.03.012> arXiv:1609.03614
- 1434 [43] Mark Last, Menahem Friedman, and Abraham Kandel. 2003. The data mining
1435 approach to automated software testing. In *Proceedings of the ninth ACM SIGKDD*
1436 *international conference on Knowledge discovery and data mining*. ACM, 388–396.
- 1437 [44] Tim Menzies. 2008. Improving IV&V Techniques Through the Analysis of Project
1438 Anomalies: Text Mining PITS issue reports-final report. *Citeseer* (2008).
- 1439 [45] Tim Menzies, Alex Dekhtyar, Justin Distefano, and Jeremy Greenwald. 2007.
1440 Problems with Precision: A Response to "Comments on 'Data Mining Static Code
1441 Attributes to Learn Defect Predictors'". *IEEE Transactions of Software Engineering*
1442 (2007).
- 1443 [46] Tim Menzies, Alex Dekhtyar, Justin Distefano, and Jeremy Greenwald. 2007.
1444 Problems with Precision: A Response to" comments on 'data mining static code
1445 attributes to learn defect predictors'". *IEEE TSE* 33, 9 (2007).
- 1446 [47] Tim Menzies, Jeremy Greenwald, and Art Frank. 2006. Data mining static code
1447 attributes to learn defect predictors. *IEEE transactions on software engineering* 33,
1448 1 (2006), 2–13.
- 1449 [48] Tim Menzies and Ying Hu. 2003. Data mining for very busy people. *Computer*
1450 36, 11 (2003), 22–29.
- 1451 [49] Tim Menzies and Andrian Marcus. 2008. Automated severity assessment of
1452 software defect reports. In *International Conference on Software Maintenance*.
1453 IEEE.
- 1454 [50] Tim Menzies, Dan Port, Zhihao Chen, and Jairus Hihn. 2005. Simple software
1455 cost analysis: safe or unsafe?. In *ACM SIGSOFT Software Engineering Notes*. ACM.
- 1456 [51] Tim Menzies and Thomas Zimmermann. 2018. Software Analytics: What's Next?
1457 *IEEE Software* 35, 5 (2018), 64–70.
- 1458 [52] Donald Michie, D. J. Spiegelhalter, C. C. Taylor, and John Campbell (Eds.). 1994.
1459 *Machine Learning, Neural and Statistical Classification*. Ellis Horwood, Upper
1460 Saddle River, NJ, USA.
- 1461 [53] Ayse Tosun Misirli et al. 2011. Ai-based software defect predictors: Applications
1462 and benefits in a case study. *AI Magazine* (2011).
- 1463 [54] Julie Moeyersoms, Enric Junqué de Fortuny, Karel Dejaeger, Bart Baesens, and
1464 David Martens. 2015. Comprehensive software fault and effort prediction: A
1465 data mining approach. *Journal of Systems and Software* 100 (2015), 80–90.
- 1466 [55] Akito Monden et al. 2013. Assessing the cost effectiveness of fault prediction in
1467 acceptance testing. *IEEE Transactions on Software Engineering* (2013).
- 1468 [56] Glenford J Myers, Corey Sandler, and Tom Badgett. 2011. *The art of software*
1469 *testing*. John Wiley & Sons.
- 1470 [57] Ashish Nadkarni and Natalya Yezhkova. 2014. Structured Versus Unstructured
1471 Data: The Balance of Power Continues to Shift. *IDC (Industry Development and*
1472 *Models) Mar* (2014).
- 1473 [58] Nicole Novielli, Daniela Girardi, and Filippo Lanubile. 2018. A Benchmark Study
1474 on Sentiment Analysis for Software Engineering Research. In *Proceedings of the*
1475 *15th International Conference on Mining Software Repositories (MSR '18)*. ACM,
1476 New York, NY, USA, 364–375. <https://doi.org/10.1145/3196398.3196403>
- 1477 [59] Alessandro Orso and Gregg Rothermel. 2014. Software testing: a research
1478 timeline (2000–2014). In *Future of Software Engineering*. ACM.
- 1479 [60] Thomas J Ostrand, Elaine J Weyuker, and Robert M Bell. 2004. Where the bugs
1480 are. In *ACM SIGSOFT Software Engineering Notes*. ACM.
- 1481 [61] Annibale Panichella, Bogdan Dit, Rocco Oliveto, Massimiliano Di Penta, Denys
1482 Poshvanyk, and Andrea De Lucia. 2013. How to Effectively Use Topic Models
1483 for Software Engineering Tasks? An Approach Based on Genetic Algorithms. In
1484 *International Conference on Software Engineering*.
- 1485 [62] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel,
1486 Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss,
1487 Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal*
1488 *of machine learning research* (2011).
- 1489 [63] Justyna Petke and Tim Menzies. 2018. Guest Editorial for the Special Section
1490 from the 9th International Symposium on Search Based Software Engineering.
1491 *Information and Software Technology* (2018).
- 1492 [64] M Porter. 1980. The Porter Stemming Algorithm. , 130–137 pages. <http://tartarus.org/martin/PorterStemmer/>
- 1493 [65] J. R. Quinlan. 1986. Induction of decision trees. *Machine Learning* 1, 1 (01 Mar
1494 1986), 81–106. <https://doi.org/10.1007/BF00116251>
- 1495 [66] Foyzur Rahman, Sameer Khatri, Earl T Barr, and Premkumar Devanbu. 2014.
1496 Comparing static bug finders and statistical prediction. In *International Conference*
1497 *on Software Engineering*. ACM.
- 1498 [67] Baishakhi Ray, Vincent Hellendoorn, Saheel Godhane, Zhaopeng Tu, Alberto
1499 Bacchelli, and Premkumar Devanbu. 2016. On the naturalness of buggy code. In
1500 *Proceedings of the 38th International Conference on Software Engineering*. ACM,
1501 428–439.
- 1502 [68] Mitch Rees-Jones, Matthew Martin, and Tim Menzies. 2017. Better Predictors for
1503 Issue Lifetime. *CoRR* arxiv.org/abs/1702.07735 (2017).
- 1504 [69] Martin Robillard, Robert Walker, and Thomas Zimmermann. 2009. Recommen-
1505 dation systems for software engineering. *IEEE software* 27, 4 (2009), 80–86.
- 1506 [70] Rainer Storn and Kenneth Price. 1997. Differential evolution—a simple and
1507 efficient heuristic for global optimization over continuous spaces. *Journal of*
1508 *global optimization* 11, 4 (1997), 341–359.
- 1509 [71] Chakkrit Tantithamthavorn, Shane McIntosh, Ahmed E Hassan, and Kenichi
1510 Matsumoto. 2016. Automated parameter optimization of classification techniques
1511

- 1497 for defect prediction models. In *2016 IEEE/ACM 38th International Conference on*
1498 *Software Engineering (ICSE)*. IEEE, 321–332.
- 1499 [72] Z. Wan, X. Xia, A. E. Hassan, D. Lo, J. Yin, and X. Yang. 2018. Perceptions,
1500 Expectations, and Challenges in Defect Prediction. *IEEE Transactions on Software*
1501 *Engineering* (2018), 1–1.
- 1502 [73] Colin White. 2005. Consolidating, Accessing and Analyzing Unstructured Data.
1503 <http://www.b-eye-network.com/view/2098>.
- 1504 [74] Leland Wilkinson, Anushka Anand, and Dang Nhon Tuan. 2011. CHIRP: A
1505 New Classifier Based on Composite Hypercubes on Iterated Random Projections.
1506 In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge*
1507 *Discovery and Data Mining (KDD '11)*. ACM, New York, NY, USA, 6–14. <https://doi.org/10.1145/2020408.2020418>
- 1508 [75] Tao Xie, Suresh Thummalapenta, David Lo, and Chao Liu. 2009. Data mining for
1509 software engineering. *Computer* 42, 8 (2009), 55–62.
- 1510 [76] Aiko Yamashita and Steve Counsell. 2013. Code smells as system-level indicators
1511 of maintainability: An empirical study. *Journal of Systems and Software* (2013).
- 1512
- 1513
- 1514
- 1515
- 1516
- 1517
- 1518
- 1519
- 1520
- 1521
- 1522
- 1523
- 1524
- 1525
- 1526
- 1527
- 1528
- 1529
- 1530
- 1531
- 1532
- 1533
- 1534
- 1535
- 1536
- 1537
- 1538
- 1539
- 1540
- 1541
- 1542
- 1543
- 1544
- 1545
- 1546
- 1547
- 1548
- 1549
- 1550
- 1551
- 1552
- 1553
- 1554
- [77] Aiko Yamashita and Leon Moonen. 2013. Exploring the impact of inter-smell relations on software maintainability: An empirical study. In *International Conference on Software Engineering*. 1555
1556
- [78] J. Yang, K. Hotta, Y. Higo, H. Igaki, and S. Kusumoto. 2012. Filtering clones for individual user based on machine learning analysis. In *International Workshop on Software Clones (IWSC)*. 1557
1558
1559
1560
- [79] Yibiao Yang, Yuming Zhou, Jinping Liu, Yangyang Zhao, Hongmin Lu, Lei Xu, Baowen Xu, and Hareton Leung. 2016. Effort-aware just-in-time defect prediction: simple unsupervised models could be better than supervised models. In *Foundations of Software Engineering*. ACM. 1561
1562
1563
1564
- [80] Shin Yoo and Mark Harman. 2012. Regression testing minimization, selection and prioritization: a survey. *Software Testing, Verification and Reliability* 22, 2 (2012), 67–120. 1565
1566
1567
- [81] Nico Zazworka, Michele A Shaw, Forrest Shull, and Carolyn Seaman. 2011. Investigating the impact of design debt on software quality. In *Workshop on Managing Technical Debt*. ACM. 1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634