

Explaining Legal Decisions Using IRAC

Trevor Bench-Capon

Department of Computer Science, University of Liverpool, Liverpool, UK.
tbc@csc.liv.ac.uk

Abstract. We suggest that the Issue, Rule, Application, Conclusion (IRAC) method can be used to produce a natural explanation of legal case outcomes. We show how a current methodology for representing knowledge of legal cases can be used to provide such explanations.

Keywords: Reasoning with cases · Explanation · Legal Reasoning

1 Introduction

Explanations given by AI and Law systems are often rather stilted and formulaic. In order to try to produce more natural explanations of legal cases we turn for inspiration to how lawyers are taught to discuss cases. IRAC stands for Issue, Rule, Application and Conclusion and is a methodology for legal analysis widely taught in US law schools¹ as a way of answering hypothetical questions posed during teaching by the Socratic method and in exams.

- The *issue* is the legal question to be answered. This is couched in terms specific to the particular case rather than in general terms. Thus can *information communicated to specific employees be considered a Trade Secret?* rather than the generic *was the information a Trade Secret?*².
- The *rule*, or rules, are the rules that are used to answer the question in the issue. For example a rule might be *information disclosed to employees is only regarded as confidential if covered by a specific non disclosure agreement*³.
- The rule must then be *applied* to the facts of the particular case being considered. For example: *The defendant was an employee of the plaintiff, and signed a general non disclosure agreement, but the particular information was not specifically mentioned.*
- The *conclusion* is the answer to the legal question. In our example: *the information is not regarded as Trade Secret since it was not covered by a specific non disclosure agreement.*

¹ For example, City University of New York (<https://www.law.cuny.edu/legal-writing/students/irac-crracc/irac-crracc-1/>) and Elizabeth Haub School of Law at Pace University (<https://academicsupport.blogs.pace.edu/2012/10/26/the-case-of-the-missing-a-in-law-school-you-cant-get-an-a-without-an-a/>).

² My examples are from the US Trade Secrets domain widely used in AI and Law [7].

³ Legal information such as this is for illustration only, and may not be an accurate reflection of the law.

Sometimes additional elements are included, such as *Rule Proof*. Rule Proof is a justification of the rule, citing the statute or case on which it is based. Thus in our example, one could cite *MBL (USA) Corp. v. Diekman* which was found for the plaintiff on the grounds that although a general employer-employee confidentiality agreement had been signed, “the court found that defendant and other employees were not told what, if anything, plaintiff considered confidential” (President Justice Downing). AI systems addressing the question of reasoning with legal cases have always attempted to explain their reasoning. Indeed the ability to provide explanations is considered (e.g. [8]) to be a major advantage of such systems over systems based on machine learning algorithms such as [13].

The key point about IRAC is that it is specifically tailored to the case at hand: it indicates what is important and different about the particular case under discussion, and uses the *specific* facts of the case to apply the general rule. This is different from the standard forms of explanation of case outcomes found in AI and Law, which go through every potential issue in the domain, obscuring the key point, and bottom out in generally applicable factors rather than specific facts. We suggest that IRAC is a more natural form of explanatory argument, just as arguments are more natural than watertight logical proofs since they can use enthymemes to suppress trivial and generally known premises and focus on the real point. In this paper we will consider how current representations can be adapted to provide IRAC style explanations.

2 Background: the US Trade Secrets Domain

US Trade Secrets has been widely used as a domain in AI and Law since its introduction in the HYPO system [15]. We will use as our primary example the US Trade Secrets case of *The Boeing Company v. Sierracin Corporation*, as modelled for the CATO system [4] and used in [1].




Fig. 1. Top level of Trade Secrets Domain from [11]

The top level of the US Trade Secrets domain is shown in Figure 1, taken from [11]. These are called *issues* in [4] and [11]. Below the issues are a number of *factors*, stereotypical patterns of fact that are legally significant and favour one or other side of the argument. In CATO [4] and many subsequent approaches,

including IBP [11] and ANGELIC [1], issues and factors are organised in a hierarchy, descending through abstract factors recognising what Alven terms “intermediate concerns” until the base level factors, the legal facts of the case, are reached. Figure 2 shows the branch of the hierarchy relating to whether the information is a Trade Secret, the left hand branch of Figure 1. CATO’s issues are a way of structuring the case, but are rather more generic than the issues which form the ‘I’ of IRAC, which are supposed to indicate what is *particular* to the case under consideration. Below we will refer to IRAC issues as *case-issues*.

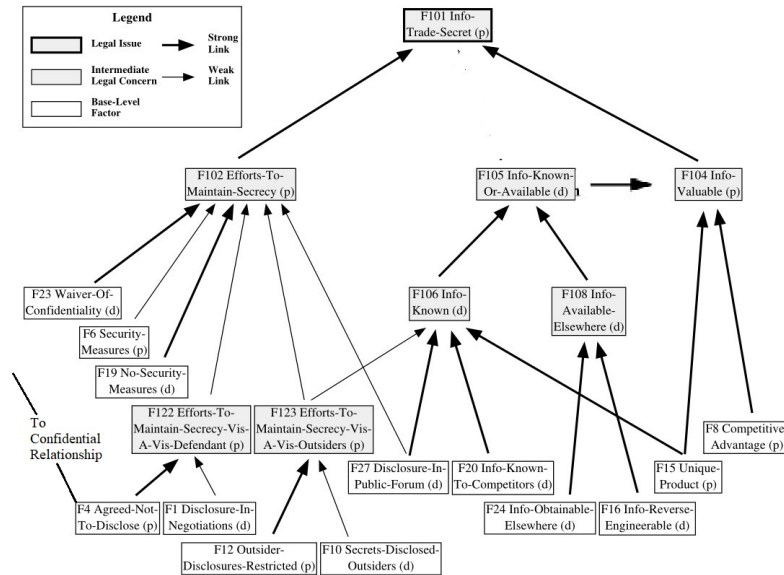


Fig. 2. Factor Hierarchy for branches relating to Info-Trade-Secret, taken from [4]

We will use a slightly simplified version of *Boeing*⁴, with six factors, five pro-plaintiff and one pro-defendant. *Boeing* was found for the plaintiff.

The factors, the side favoured and their associated CATO issues are:

- F4 NonDisclosureAgreement (plaintiff); Maintain-Secrecy, ConfidentialRelationship. This also appears in the branch not shown in Fig 2.
- F6 SecurityMeasures (plaintiff): Maintain-Secrecy
- F10 SecretsDisclosedOutsiders (defendant): Maintain-Secrecy
- F12 OutsiderDisclosuresRestricted (plaintiff): Maintain-Secrecy
- F14 RestrictedMaterialsUsed (plaintiff): Improper-Means (not in Fig 2)
- F21 KnewInfoConfidential (plaintiff): Confidential-Relationship

⁴ The simplification is that we do not use F1, Disclosure-in Negotiations, since we consider it subsumed by F10, since the defendant is an outsider.

The F numbers are the identifiers given in [4], shown in Figure 2, and used in subsequent discussions such as [1] and the Tables below.

Here it should be noted that since improper (although not criminal) means were used to obtain the information, by improperly exploiting restricted materials, there is no question but that the information should be regarded as misappropriated. So even though there is some support for the other arm, ConfidentialRelationship, it is not needed: ImproperMeans alone is enough for misappropriation. So the issue is whether information disclosed to outsiders can be considered a trade secret. Here the answer is that it can be, since the disclosures were subject to restrictions, showing that the information was considered secret by the plaintiff, and adequate efforts had been taken to maintain the secrecy. The value of the information is not disputed, and so can be assumed by default.

3 Standard explanation in AI and Law

One approach to building an AI and Law system is to formalise the relevant legal knowledge and to elicit the facts of the particular case by querying the user. This is the approach of the classic [17], which formalised the British Nationality Act. The explanation was the standard *how?* typical of expert systems of the time. This approach remains relevant today, as shown by the ANGELIC methodology [1], which has recently been used to build a fielded system [3]. In [1] the Boeing case was used as an example. The proof trace of the Prolog program was post-processed to give the following explanation (note that the ANGELIC program uses defaults to resolve issues for which there are no base level factors):

We find for plaintiff. The information was a trade secret: efforts were taken to maintain secrecy, since disclosures to outsiders were restricted and the defendant entered into a non-disclosure agreement and other security measures were applied. The information was unique. It is accepted that the information was valuable and it is accepted that the information was neither known nor available. A trade secret was misappropriated: there was a confidential relationship since the defendant entered into a non-disclosure agreement and it is accepted that the information was used. Moreover improper means were used since the defendant used restricted materials.

This follows the usual pattern of a rule based *how?* explanation: each component of the and/or tree of Figure 1 is considered, and the result justified by citing lower level rules until the base level factors (given or accepted by default) are reached. The explanation has a conclusion and a rule (here implicit, although older programs such as [17] cite the rule explicitly). What is missing here is the focus provided by the case-issue: the explanation covers all the elements without any effort to identify the important point, or to say why the factors are present.

Case based reasoning in AI and Law also typically considers the case as a whole. What would happen in HYPO is that the case base would be searched until the closest match in terms of shared factors was found. A possible match,

taken from the limited set of publicly available cases listed in [12], is *Trandes Corp. v. Guy F. Atkinson Co.* *Trandes* had just three factors, all shared with *Boeing*: *AgreedNotToDisclose*, *SecurityMeasures* and *SecretsDisclosedOutsiders* (again we disregard F1 as subsumed by F10). It was found for the plaintiff. A case based explanation (based on the explanation of a different case in [16]) would cite the factors in common:

Where the plaintiff had taken security measures, the defendant had agreed not to disclose the information and information had been disclosed to outsiders, Plaintiff should win claim. Cite *Trandes*.

The defendant might now reply: *In Trades the disclosures were of a lesser extent than Boeing*. This distinction cannot be made with Boolean factors as used in CATO, IBP and [1], but is available in [9], where factors can have magnitudes.

For rebuttal the plaintiff can say that in *Boeing* the disclosures to outsiders were restricted: this is not so in *Trandes*, which makes *Boeing* significantly stronger than *Trandes*. This style of factor based explanation remains relevant, and is currently advocated as a means of explaining systems based on Machine Learning techniques [10]

It might now be asked why *Trandes* was found in favour of the plaintiff when the information had been disclosed to outsiders without restriction. From a reading of the decision it can be seen that these disclosures were held to be too limited to be of significance: “Although *Trandes* did give WMATA a single demonstration disk in contemplation of a licensing agreement, it did so only in confidence. Such limited disclosure does not destroy secrecy” (opinion delivered by Williams, Circuit Judge). This may cast doubt on whether this factor should have been ascribed to *Trandes* at all. The issue in *Trandes* was really whether the disclosures had been sufficient to compromise the secret, whereas in *Boeing* they clearly were. Because the case based explanation uses the whole case it may rely on similarities which were not relevant to the crucial issue. Again, what is missing is the focus the case-issue provides.

A key point about these explanations is that they consider the case as a whole. Thus the rule based explanation goes through all the CATO issues without distinguishing on which of them the case turns. The case based explanation considers all factors in common with the precedent. Neither focus on a specific issue, particular to the case, which is what we want for our IRAC case-issue.

A second point is that we can see two types of case-issue. In *Boeing* we have two factors favouring different sides in the same branch of the tree: *SecretsDisclosedOutsiders* and *OutsiderDisclosuresRestricted*, bringing into question whether or not their parent factor, *EffortsMadeToMaintainSecrecyVis-a-VisOutsiders*, is present. If it is not, then, as can be seen in Figure 2, we can decide that the efforts to maintain secrecy were not adequate and so the information cannot be regarded as a trade secret. We term this a *conflict-issue*. In *Trandes*, in contrast, the issue is whether or not we consider that information was disclosed to outsiders to a sufficient extent: that is whether we should consider this factor to be meaningfully present in the case. We will term this an *ascription-issue*.

4 Issues in AI and Law

As mentioned above, CATO [4] used *issue* to describe the top levels of its abstract factor hierarchy. This practice was also used in Issue Based Prediction (IBP) [11], [5] and the ANGELIC methodology [1], [2]. CATO was concerned with helping law students to distinguish a case, and so does not give an explanation of the outcome. It does, however, use issues to organise its advice. IBP, however, was directed towards predicting an outcome. It first identified the relevant issues, those with at least one factor present in the case in its sub-tree, and where there were opposing factors relevant to an issue, used matching precedents to choose a resolution. The explanation provided, like that from ANGELIC given above, proceeds on an issue by issue basis, rather than identifying and focussing on the crucial *case-issue*, but at least uncontested issues are ignored.

In the remainder of the paper I will discuss how an IRAC explanation can be produced from a system constructed using the ANGELIC methodology.

5 CATO in ANGELIC

The ANGELIC methodology produces an ADF corresponding to the factor hierarchy of [4], part of which is shown in Fig 2. Each node of the ADF can have positive or negative children. The ADF for the issues are shown in Table 1. The ADF for the abstract factors are shown in Table 2. Factors present or inferred in *Boeing* are in bold, those in *Trandes* in italics. Base Level factors are given using CATO F numbers [4] and Fig 2. The value and use of in the information was not contested, and so no related factors were mentioned in either case.

CATO ID	Name	Positive Children	Negative Children
F200	TradeSecretMisappropriation	<i>F201</i> , F203	F124
F203	InfoTradeSecret	<i>F102</i> , F104	
F104	InfoValuable	F8, F15	F105
F102	EffortstoMaintainSecrecy	<i>F6</i> , <i>F122</i> , <i>F123</i>	F19, F23, F27
F201	InfoMisappropriated	<i>F110</i> , F112, <i>F114</i>	
F112	InfoUsed	F7, F8, F18	F17
F114	ConfidentialRelationship	<i>F115</i> , <i>F121</i>	
F110	ImproperMeans	<i>F111</i>	F120

Table 1. IBP Logical Model as an ADF taken from [1]. Factors present in *Boeing* are in bold. Factors in *Trandes* are in italics

6 Modelling IRAC with ANGELIC

As we saw above, issues can be of two types: conflict based and ascription based. We will consider each in turn.

CATO ID	Name	Positive Children	Negative Children
F105	InfoKnownOrAvailable	F106, F108	
F106	InfoKnown	F20, F27	F15, F123
F108	InfoAvailableElsewhere	F16, F24	
F111	QuestionableMeans	F2, F14 , F22, F26	F1, F17, F25
F115	NoticeOfConfidentiality	F4, F13, F14, F21	F23
F120	LegitimatelyObtainable	F105	F111
F121	ConfidentialityAgreement	<i>F4</i>	F5, F23
F122	MaintainSecrecyDefendant	<i>F4</i>	<i>F1</i>
F123	MaintainSecrecyOutsiders	F12	F10
F124	DefendantOwnershipRights	F3	

Table 2. CATO factors as ADF taken from [1]. Factors present in Boeing are in bold. Factors in *Trandes* are in italics

6.1 Conflict Issues

Examination of Table 1 shows that there are no case-issues at the CATO issue level, because there are no negative children. Two leaf issues in Fig 1 are relevant, *EffortstoMaintainSecrecy* and *InfoMisappropriated* (*ConfidentialRelationship* is only relevant when *InfoUsed* is also relevant and *InfoValuabe* is uncontested), but since both have only positive factors, the case is clearly decidable for the plaintiff at this level. So, to find the case-issue we must delve deeper down the tree, and examine Table 2.

In Table 2 we see that there are is one factor with conflicting base level factors, *MaintainSecrecyOutsiders*. This conflict will thus be our case-issue. *MaintainSecrecyOutsiders*, if resolved differently, could have destroyed the plaintiff’s claim by removing *EffortstoMaintainSecrecy*, and hence showing the information not to be regarded as a Trade Secret. This then is the case-issue on which the case turns. Since any case-issue requires a factor with a positive child and a negative child, we can now state a case-issue using a template of the form *Can the plaintiff be considered to Parent when Negative Child given that Positive Child?*:

- Can the plaintiff be considered to *MaintainSecrecyOutsiders* when *SecretsDisclosedOutsiders* given that *OutsiderDisclosuresRestricted*?

We now examine the acceptance conditions for the node F123 at which the conflict occurs. We take the acceptance conditions from [1]:

```
ACCEPT IF OutsiderDisclosuresRestricted
REJECT IF SecretsDisclosedOutsiders
ACCEPT
```

The default ACCEPT indicates that the burden of proof for this factor is on the defendant, and the order of the conditions indicate the priority of the two rules. The acceptance conditions can be annotated with the case or cases which led to their inclusion. These can then be used as the Rule Proof, if that element is required. The rule of the case is thus

– MaintainSecrecyOutsiders if OutsiderDisclosuresRestricted

The application is that OutsiderDisclosuresRestricted is present in the case, and so the antecedent is satisfied. This could be elaborated using specific facts taken from the decision “Boeing did not lose its secrets through confidential disclosure of drawings to Libbey. The secrets were preserved by first Libbey’s and then Sierracin’s promise to keep the information confidential.” (opinion delivered by Justice Dore). We suggest an extension of the ANGELIC methodology so that that when a case is analysed into factors, the extract from the opinion on which the ascription was based is recorded, so that it can be used in the explanation.

The conclusion is that Boeing did MaintainSecrecyOutsiders, and hence made adequate efforts to maintain secrecy, and that they had a Trade Secret which was misappropriated through the use of restricted materials. We now have all the elements of IRAC, which can be used as a natural explanation, focussing on what is important in the particular case.

In Boeing the issue is whether the plaintiff can be considered to have maintained secrecy with respect to outsiders when the plaintiff disclosed the information to outsiders when these outsider disclosures were restricted? We apply the rule: if outsider disclosures are restricted then secrecy with respect to outsiders is maintained. In *Boeing*, the secrets were preserved by first Libbey’s and then Sierracin’s promise to keep the information confidential. Thus secrecy with respect to outsiders was maintained, and the information can be regarded as a trade secret.

6.2 Ascription Based Issues

We now consider *Trandes*. In *Trandes* there were no confidentiality restrictions on the disclosures to outsiders, and so it would seem a straightforward case in which the information is not a trade secret because it was in the Public Domain. This, however, is only the case if we insist that factors are Boolean. This is true of CATO, but not of HYPO, which used dimensions indicating the extent to which a particular feature favoured a party. Recently, the recognition that factors are present to differing extents is increasingly recognised (e.g. [14]). The ANGELIC methodology has been extended to accommodate factors present to differing extents as reported in [9]. In [9] this was applied to CATO, and one of the factors treated as non-Boolean was SecretsDisclosedOutsiders. In Williams’ decision in *Trandes* we can find the form of the alleged disclosure:

The advertisement described in general terms the capabilities of the Tunnel System and offered a demonstration disk for \$100. The demonstration disk did in fact contain an object-code version of the Tunnel System, but was configured to operate in demonstration mode only.

Note that the users did not have open access to the source code constituting the Trade Secret. So the disclosure would seem to be minimal: moreover the advertisement attracted very few enquiries. Thus for *Trandes*, we may include

SecretsDisclosedOutsiders (as argued by the defendant), but the extent will be small and so may fail to meet the threshold needed to defeat MaintainSecrecy-Outsiders. Thus the information can continue to be seen as a trade secret.

The issue here is thus whether distribution of a demonstration disk containing object level code configured to operate in demonstration mode only means that inadequate security measures were taken. The rule that was applied by Williams was “limited disclosure does not destroy secrecy”, with *Space Aero v Darling* cited as the precedent. We now apply the rule by stating that the disclosure in *Trades* was sufficiently limited.

Thus the ascription issue can be identified by looking for a non boolean factor, in particular one below or close to the threshold at which it will influence the acceptability of its parent. In *Trandes* the factor SecurityMeasures also has magnitude, but these were quite extensive, and so the presence of this factor is not in dispute. The conclusion is that *Trandes* did MaintainSecrecyOutsiders, and hence made adequate efforts to maintain secrecy, and that they had a Trade Secret which was misappropriated through the use of information disclosed to the defendant in a relation of confidence.

In *Trandes* the application of the rule that can be extracted from the representation is simply the claim that the disclosures were too limited to destroy secrecy, backed up the quotation showing the nature of the disclosure. We may, however, encounter rather more extensive discussion of whether the factor should be applied or not. This may require true analogical reasoning, as discussed in [18]. Her example was a hypothetical based on *Dillon v. Legg* in which the issue is whether a kindergarten teacher who witnessed an accident involving one of her pupils can be considered sufficiently analogous to a mother (the relationship in *Dillon*) to receive damages for emotional distress. On the one hand both are in a close caring relationship with the child, but on the other there is no blood relationship: the precedent in question simply states that “how closely the plaintiff was related to the victim” is what needs to be considered. Here we have a HYPO style dimension and need to decide where to fix the point at which it switches from pro-plaintiff to pro-defendant [14]. In [6] it was suggested that such reasoning requires too much common sense knowledge of the world to be achievable in current AI and Law systems. A rich ontology might be produced for past cases and so used for teaching, but not one broad enough to cover arbitrary future cases, as would be required for prediction.

Thus, for ascription issues, we must be content with the limited notion of application provided by thresholds. If a good discussion of application is, as the second url in footnote 1 suggests, required for an A grade, we must perhaps be content with a B minus.

7 Conclusion

In this paper we have proposed the IRAC method as a natural way of explaining legal cases, which focusses on the central point on which the case turns. We showed how this might be produced from a factor based representation such is

produced by the ANGELIC methodology [1]. We further identified two types of issue: *conflict issues*, which turn on opposing factors relating to a common parent, and *ascription issues*, turning on whether a factor is present to a sufficient extent. The representation allows us to deal satisfactorily with conflict issues, but even if we can represent the extent to which factors are present in a case our explanation of the application of rules relating to ascription issues will be limited in cases where true analogical reasoning is needed.

References

1. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
2. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Statement types in legal argument. In *Proceedings of JURIX 2016*, pages 3–12, 2016.
3. L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, S. Whittle, R. Williams, and C. Wolfenden. Noise induced hearing loss: Building an application using the ANGELIC methodology. *Argument and Computation*, 10(1):5–22, 2019.
4. V. Alevén. *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh, 1997.
5. K. D. Ashley and S. Brüninghaus. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*, 17(2):125–165, 2009.
6. K. Atkinson and T. Bench-Capon. Reasoning with legal cases: Analogy or rule application? In *Proceedings of the Seventeenth ICAIL*, pages 12–21, 2019.
7. T. Bench-Capon. HYPO’s legacy: Introduction to the virtual special issue. *Artificial Intelligence and Law*, 25(2):1–46, 2017.
8. T. Bench-Capon. The need for good old fashioned AI and Law. In W. Hötzenndorfer, C. Tschohl, and F. Kummer, editors, *International trends in legal informatics: a Festschrift for Erich Schweighofer*, pages 23–36. Weblaw, Bern, 2020.
9. T. Bench-Capon and K. Atkinson. Lessons from implementing factors with magnitude. In *Proceedings of JURIX 2018*, pages 11–20, 2018.
10. L. K. Branting, C. Pfeifer, B. Brown, L. Ferro, J. Aberdeen, B. Weiss, M. Pfaff, and B. Liao. Scalable and explainable legal prediction. *AI and Law*, Online 2020.
11. S. Brüninghaus and K. D. Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th ICAIL*, pages 233–242. ACM, 2003.
12. A. Chorley and T. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *Artificial Intelligence and Law*, 13(3-4):323–371, 2005.
13. M. Medvedeva, M. Vols, and M. Wieling. Using machine learning to predict decisions of the European Court of Human Rights. *AI and Law*, 28:237–266, 2020.
14. A. Rigoni. Representing dimensions within the reason model of precedent. *Artificial Intelligence and Law*, 26(1):1–22, 2018.
15. E. L. Rissland and K. D. Ashley. A case-based system for Trade Secrets law. In *Proceedings of the 1st ICAIL*, pages 60–66. ACM, 1987.
16. E. L. Rissland and K. D. Ashley. A note on dimensions and factors. *Artificial Intelligence and Law*, 10(1-3):65–77, 2002.
17. M. Sergot, F. Sadri, R. Kowalski, F. Kriwaczek, P. Hammond, and H. Cory. The British Nationality Act as a logic program. *Comm. ACM*, 29(5):370–386, 1986.
18. K. Stevens. Reasoning by precedent—between rules and analogies. *Legal Theory*, pages 1–39, 2018.