Legal Knowledge and Information Systems
G. Sileno et al. (Eds.)

doi:10.3233/FAIA230941

The Importance of Intermediate Factors

Ilaria CANAVOTTO [a,1] and John HORTY [b,2]
[a] *Philosophy Department, University of Maryland, College Park, USA*
[b] *Philosophy Department and Institute for Advanced Computer Studies, University of Maryland, College Park, USA*

Abstract. Bench-Capon argues that intermediate factors have no role to play in precedential constraint. We offer a constrasting perspective.

Keywords. Reason model, hierarchies, issues, intermediate factors.

1. Introduction

In previous work [1], we developed a version of the reason model of precedential constraint applicable in a rich hierarchical setting, with a variety of intermediate factors lying between base-level factors and high-level issues. Although the theory set out there is a conservative extension of the standard reason model [2], we also suggested that the presence of intermediate factors makes an important difference. But what kind of difference? A reductive suggestion might be that the importance of intermediate factors is merely cognitive, rather than logical. Intermediate factors may help a court to structure its reasoning as it moves from the base-level factors present in a case to an overall judgment, but do not affect the meaning of the judgment itself: once a judgment has been reached, the intermediate factors that guided the court's reasoning can be forgotten and the chain of intermediate decisions justifying that outcome compressed into a single step.[3]

One goal of our earlier paper was to show that this reductive view is wrong: decisions about intermediate factors have logical, not just psychological, significance. In order to establish this result, we defined a notion of flattening that maps rich hierarchies, cases, and case bases into corresponding flattened structures, and then showed that the reductive suggestion fails in two directions. First, constraint in a rich hierarchical setting does not entail constraint in the corresponding flattened setting. This initial observation was not surprising. Everyone knows that certain patterns are apparent only from a higher-level perspective, an idea that has been exploited across the field of AI, from chunking in the SOAR architecture [4] to hierarchical task network planning [5]. Second, and more surprisingly, we also showed that new relations of constraint might appear among flattened structures that were not apparent in the original hierarchical setting.

We understood our first observation—that the reason model of constraint can be applied to a variety of intermediate factors—as a generalization of Bench-Capon and Atkinson's [6] suggestion that the reason model of precedential constraint should be applied at the level of issues, not just ultimate outcomes. In a recent paper, however, Bench-Capon [7] has argued that only issues, and not intermediate factors, should affect

[1] E-mail: icanavot@umd.edu.
[2] E-mail: horty@umd.edu.
[3] This view mirrors Goodhart's proposal [3] that the meaning, or *ratio decidendi*, of a case is exhausted by the connection between base-level factors and the ultimate judgment.

the notion of constraint in a real, substantive way; he also suggests some problems with the presentation of our second result, that new relations of constraint can appear among flattened structures.

Bench-Capon raises important issues and advances the discussion in a useful way. The purpose of this note is to provide our perspective. We begin in the next section by reviewing and attempting to standardize terminology, summarizing our notion of flattening, and then presenting some examples to establish the difference between hierarchical and flat constraint. We then respond to Bench-Capon's criticisms.

2. Background

2.1. Terminology

Factor hierarchies contain two kinds of factors: *base-level* factors and *abstract* factors. Base-level factors are concepts about whose application there is no disagreement, while abstract factors are concepts about whose application there may be some dispute. Abstract factors can themselves be divided into *intermediate* and *top-level* factors. Top-level factors—or *issues*—are abstract factors that are so high in a hierarchy that any further relation between these concepts and the case outcome is purely deductive, while *intermediate* factors are abstract factors that are neither so high in the hierarchy that their relation to an overall judgment is simply deductive nor so low in the hierarchy that questions concerning their application can be resolved entirely without dispute. The ultimate outcome of a case, a decision for the plaintiff (π) or defendant (δ), can itself be classified as an issue—it is the topmost issue of any hierarchy to which it belongs.[4]

Two well-known examples of factor hierarchies are found in CATO [8] and IBP [9]. The CATO hierarchy contains a set of base-level factors, multiple layers of intermediate factors, and a set of issues that does not include the ultimate outcome. Factor-based reasoning moves step-by-step from the base-level factors through the different layers of intermediate factors, until a resolution of the issues at the top of the hierarchy is reached. The IBP hierarchy, also appealed to in VJAP [10], likewise contains base-level factors and a set of issues—which, this time, also includes the outcome. Unlike the CATO hierarchy, the IBP hierarchy does not contain any intermediate factors. Accordingly, factor-based reasoning proceeds, in a single step, from the base-level factors to a resolution of the issues on which they bear. Logical reasoning is then used to deduce the ultimate outcome from decisions on these issues.

The reason model of constraint was originally formulated [2] in the context of a *flat* hierarchy, containing only base-level factors directly supporting the ultimate outcome, a decision for π or δ. The notion of constraint defined there can thus be characterized as *flat* constraint, or *F-constraint*. More recently, Bench-Capon and Atkinson argued that the reason model should be formulated against the background of an IBP-style hierarchy, with base-level factors directly supporting, not ultimate outcomes, but rather, decisions concerning a number of issues bearing on that outcome. On their view, instead of constraining an overall judgment, factor-based decisions are better understood as constraining the resolution of these issues—a decision for π or δ is then supposed to follow from decisions on these issues through ordinary logic.

Interestingly, before precise models of precedential constraint appeared in AI and Law, both Branting [11] and Prakken and Sartor [12] made a similar, but slightly more general suggestion, allowing constraint to derive from previous decisions concerning, not only issues, but also intermediate factors. Our goal in [1] was to develop this suggestion

[4] Bench-Capon [7] adopts a slightly different terminology, using the term "factor" to refer to what we call a base-level factor and "abstract factor" to what we call an intermediate factor.

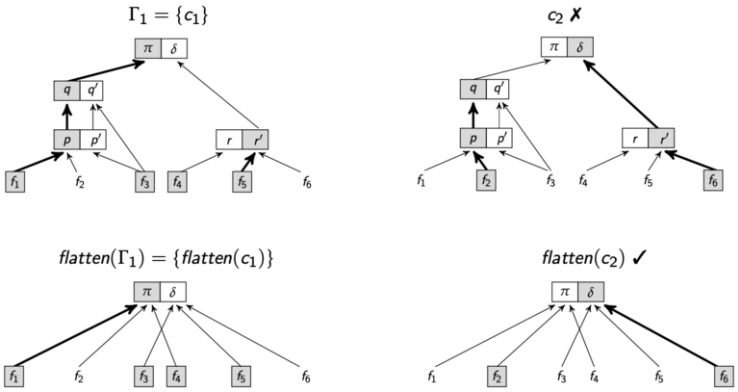

Figure 1. H-constraint without F-constraint

within the reason model of constraint. We offered a precise formalization of the idea that decisions concerning issues might depend, in a step-by-step fashion, on decisions concerning a number of intermediate factors and that constraint would then derive from these intermediate decisions as well; this new notion of constraint can be described as *hierarchical* constraint, or *H-constraint*.[5]

2.2. Flattening

The question now arises: What is the relation between the original notion of F-constraint and the new notion of H-constraint? To investigate this question, we introduced the concept of *flattening*, defined formally in [1], but which we can summarize here. The flattening of a case c decided in the context of a factor hierarchy $\mathscr{H}$ involves three steps. First, the factor hierarchy $\mathscr{H}$ is transformed into the flat hierarchy $flatten(\mathscr{H})$: all intermediate factors are removed, and every base-level factor is linked directly to the ultimate outcome it favors.[6] Next, the hierarchical justification j from the original case is likewise transformed into the flat justification $flatten(j)$: the reason for the decision concerning an issue is projected onto the set of base-level factors favoring components of that reason. Finally, the new case $flatten(c)$ from the flattened hierarchy $flatten(\mathscr{H})$ is defined as containing the fact situation from the original case c, the flattening $flatten(j)$ of the original hierarchical justification j from c, and the original outcome of c.

To illustrate these ideas, begin with the case c_1 from the top left of Figure 1. This case is decided in the context of a factor hierarchy—say, $\mathscr{H}$—that contains six base-level factors (f_1 to f_6), six intermediate factors (p, q, r and their contraries p', q', r'), and the two top-level issues π and δ. This case represents a scenario in which the court was confronted with the fact situation $X_1 = \{f_1, f_3, f_4, f_5\}$ and found for π with the following hierarchical justification j_1, depicted with thick arrows: *p holds because of* f_1 *and* r' *holds because of* f_5. *In addition, q holds because of p. In spite of the fact that* r' *supports* δ, *we decide for* π *because of q.*

The case c_1 can now be flattened following our three steps. First, the underlying hierarchy $\mathscr{H}$ is transformed into the flat hierarchy $flatten(\mathscr{H})$, shown at the bottom left

[5] An alternative model also based on the idea that decisions about intermediate factors may constrain future courts was presented in [13].

[6] The notion of favoring must itself be provided with a formal definition—but very roughly, a base-level factor f favors an abstract factor s in the hierarchy if the hierarchy contains a sequence of positive links leading from f to s.

of Figure 1. Next, the hierarchical justification j_1 is itself flattened by projecting q, the only reason in j_1 that favors an issue, onto f_1, the only base-level factor from X_1 that favors q. The result is *flatten*(j_1): *In spite of the fact that* f_3 *and* f_4 *support* δ*, we decide for* π *because of* f_1. Finally, the new case *flatten*(c_1), depicted on the bottom left of the figure, consists of the original fact situation X_1, the flat justification *flatten*(j_1), and the original outcome π.

2.3. Examples

This notion of flattening can be extended to a case base in the obvious way, by defining the flattening of a case base Γ as the set of flattenings of the cases from that case base—that is, by taking *flatten*$(\Gamma) = \{$*flatten*$(c) : c \in \Gamma\}$. And then our original question concerning the relation of F-constraint to H-constraint can be refined. Given a case base Γ developed in the hierarchy $\mathscr{H}$, where *flatten*(Γ) is the flattening of this case base in the context of *flatten*$(\mathscr{H})$, we can now ask: (1) If Γ requires a particular decision by H-constraint in the new situation X, does it follow that *flatten*(Γ) requires that same decision by F-constraint? And (2): If *flatten*(Γ) requires a particular decision in X by F-constraint, does Γ require that same decision by H-constraint?

The answer to question (1) is negative—there are decisions that are required by a case base by H-constraint that are not required by the flattening of that case base by F-constraint. Figure 1 provides an example. As we have seen, c_1, in the top left, represents a scenario in which $X_1 = \{f_1, f_3, f_4, f_5\}$ was decided for π in the context of the hierarchy $\mathscr{H}$ on the basis of the hierarchical justification j_1. As in the standard reason model, then, c_1 induces a priority ordering among reasons: the four decisions contained in the justification j_1 tell us that f_1 is more important than f_3, that f_5 is more important than f_4, that p is more important than the empty reason, and that q is more important than r'. The reason model requires later courts to respect the priorities established by earlier courts. Accordingly, against the background of $\Gamma_1 = \{c_1\}$, a future court would not be permitted to decide the new fact situation $X_2 = \{f_2, f_6\}$ for δ, as in the case c_2, on the basis of the hierarchical justification j_2 highlighted in the the top right of the figure. Why not? Because the c_2 decision for δ because of r' is inconsistent with the c_1 decision for π because of q—according to the c_2 decision, r' is more important than q, but it was already established in c_1 that q is more important than r'. Since j_2 is the only hierarchical justification based on X_2 in the context of $\mathscr{H}$ that supports δ, we conclude that Γ_1 requires the court to find X_2 for π by H-constraint.

But does the flattened case base *flatten*(Γ_1) developed in the context of the flat hierarchy *flatten*$(\mathscr{H})$ likewise require X_2 to be decided for π by F-constraint? No. To see this, observe that the case *flatten*(c_1), in the bottom left of the figure, tells us only that f_1 is more important than the combination of f_3 and f_5. Accordingly, against the background of *flatten*(Γ_1), a future court would be permitted to find X_2 for δ, as in the decision *flatten*(c_2), on the basis of the *flatten*(c_2) justification *flatten*(j_2) highlighted in the the bottom right—the *flatten*(c_2) decision that δ holds because of f_6 tells us only that f_6 is more important than f_2, which is perfectly consistent with the *flatten*(c_1) decision that f_1 is more important than f_3 and f_5. Therefore, although Γ_1 requires X_2 to be decided for π by H-constraint, *flatten*(Γ_2) does not require the same decision by F-constraint.

What about question (2)? If *flatten*(Γ) requires a particular decision in X by F-constraint, does Γ require that same decision by H-constraint? Surprisingly the answer to this question is also negative. Figure 2 provides an example. In c_3, top left, the court decided $X_3 = \{f_1, f_4, f_5\}$ for π on the basis of the hierarchical justification j_3: *p holds because of* f_1 *and r holds because of* f_4*. In addition, q holds because of p. So we find for* π *because of q*. It is easy to see that, against the background of $\Gamma_3 = \{c_3\}$, a future court would be permitted to decide $X_4 = \{f_1, f_5\}$ in favor of δ, as in c_4, on the basis of the justification j_4 highlighted in the hierarchy on the top right—there is no inconsis-

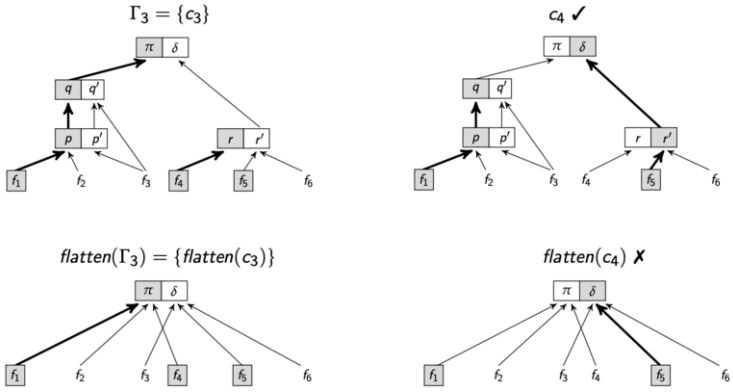

Figure 2. F-constraint without H-constraint

tency between the priorities among reasons derived from c_4 and those derived from the precedent case c_3. Hence, Γ_3 does not require the court to find X_4 for δ by H-constraint.

In this case, however, $flatten(\Gamma_3)$ does require X_4 to be decided for π. To see this, note that $flatten(c_3)$, in the bottom left of the figure, tells us that f_1 is more important than f_5. Accordingly, against the background of $flatten(\Gamma_3)$, a future court would not be permitted to find X_4 for δ, as in $flatten(c_4)$, on the basis of the flattened justification $flatten(j_4)$, highlighted in the hierarchy on the bottom right: the decision that δ holds because of f_5 would tell us that f_5 is more important than f_1, which is inconsistent with the $flatten(c_3)$ judgment that f_1 is more important than f_5. Since $flatten(j_4)$ is the only justification in the flattened hierarchy based on X_2 that supports δ, we can conclude that $flatten(\Gamma_3)$ requires the court to find X_4 for π by F-constraint.

3. Discussion

Since our analysis allows both H-constraint without F-constraint and F-constraint without H-constraint, there are, as Bench-Capon notes, four possibilities: given a background case base, a new decision might be both H-constrained and F-constrained, neither H-constrained nor F-constrained, H-constrained but not F-constrained, or F-constrained but not H-constrained. The first two possibilities, where the two notions of constraint coincide, pose no real problem. But in the third and fourth possibilities, where H-constraint and F-constraint differ, we must ask: Which is correct—which of the two formal notions represents the real concept of precedential constraint?

3.1. H-constraint without F-constraint

We suggested in [1] that H-constraint was real, and important to recognize. Bench-Capon, however, argues that F-constraint is the correct notion—so that, when a new situation is H-constrained but not F-constrained, it is not really constrained. Only F-constraint, he thinks, is real constraint.

Bench-Capon's primary reason for this conclusion is that H-constraint relies on intermediate factors, which are not a legitimate part of the law. In his view, top-level abstract factors—that is, issues—are indeed part of the law:

> Issues are typically found in the legislation ... or framework precedents ... or authoritative commentaries. The law is expressed in terms of issues, and judges are obliged

> to consider and resolve the issues and decide the case by using the and/or tree of issues.[7]

And it is precisely because issues are part of the law that previous decisions concerning issues have precedential force, so that the issue-based reason model [6] is itself legitimate. Base-level factors are likewise meaningful. They may "correspond to commonly used phrases in legal decisions" or simply "commonly considered aspect of the cases." But when it comes to intermediate factors—abstract factors from a hierarchy other than the top-level issues—Bench-Capon argues that these

> ... do not form part of the law. The factor hierarchies are products of the analysis performed by *knowledge engineers*. They are intended to provide a *descriptive* model of how judges have resolved the issues in past cases. There is *no suggestion* that judges use these hierarchies, or would acknowledge them. They are simply there as part of an effort to model the reasoning of judges.

And since the intermediate factors used to define the notion of H-constraint have doubtful legal status, he concludes that the notion of H-constraint is itself illegitimate.

To buttress his conclusion, Bench-Capon also notes that intermediate factors were originally introduced in CATO, not as real legal entities in their own right, but only "to group together [base-level] factors that could be considered to substitute for or cancel one another, which was a particular requirement of CATO's motivation," and then speculates that their "lack of legal status" led to their omission from later work in the same tradition, such as IBP and VJAP.

We do not disagree with Bench-Capon's analysis of the particular legal domains studied in CATO, IBP, or VJAP—he may well be right that the intermediate factors identified in CATO have little real legal meaning, and are omitted from IBP and VJAP for exactly that reason. We are not experts in the domain, and cannot question this judgment.

Rather than focusing on a particular legal domain, however, we take a broader perspective. Why do we study AI and Law? One reason is simply to bring computational techniques to bear on the study of law itself, with the goal of developing tools useful in, for example, teaching legal argumentation or predicting case outcomes. Our reasons are different. We want to understand normative reasoning more generally and believe the ideas developed in AI and Law have much wider applicability than the analysis of the law per se. They should be useful in the field of Ethics and AI, for example, or in Explainable AI.[8] For these reasons—because we are concerned with broader applicability, and also because we are not lawyers—we tend to motivate our analyses with a broader range of examples, including artificial examples.[9]

With this in mind, let us consider a concrete interpretation of the example from Figure 1 to see exactly why, in this case, H-constraint without F-constraint is plausible. Suppose Jack and Jo are the parents of two children, Emma and Max, who occasionally ask for treats. In deciding whether to grant these requests, Jack and Jo rely on the earlier hierarchy $\mathscr{H}$, with factor meaning now indicated by the labels in Figure 3. They agree, that is, that: Whether a child can *have ice cream* (π) depends on whether that child *behaved at home* (q) and *misbehaved at school* (r'). Whether a child behaved at home depends on whether that child *tidied up their room* (p), which itself depends on whether that child *folded their clothes* (f_1), *made their bed* (f_2), and *threw toys on the floor* (f_3). And that whether a child misbehaved at school depends on whether that child *turned in*

[7] All quotations in this paragraph and the next are from Section 4 of [7].

[8] For an application of ideas from AI and Law in Ethics and AI, see [14]; for application in Explainable AI, see [15,16,17].

[9] For other uses of artificial examples to illustrate legal concepts, see [18,19].

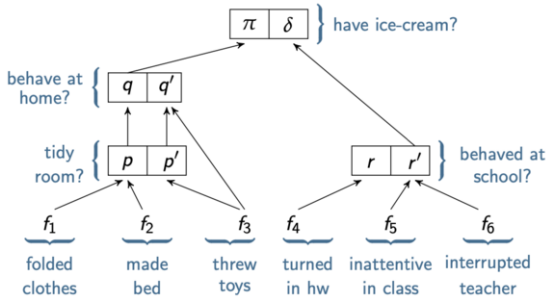

Figure 3. A concrete interpretation of the hierarchy $\mathscr{H}$

their homework (f_4), was *inattentive in class* (f_5), and *interrupted their teacher* (f_6).[10]

Given this interpretation of the factor hierarchy, suppose that, one day, Emma, who folded her clothes (f_1) and turned in her homework (f_4), but threw her toys on the floor (f_3) and was inattentive in class, (f_5), asks Jo for ice cream; this situation can be represented through the previous $X_1 = \{f_1, f_3, f_4, f_5\}$. And the case c_1 can then be taken to represent Jo's decision to let Emma have ice cream, where the earlier justification j_1 from that case can be interpreted as follows: *Emma tidied up her room because she folded her clothes, while she misbehaved at school because she was inattentive in class. Since she tidied up her room, Emma behaved at home. In spite of the fact that Emma's misbehavior at school is a reason against letting her have ice cream, my decision is that Emma can have ice cream because she behaved at home*.

Now, suppose that, after Jo's decision concerning Emma, Max, who made the bed (f_2) but interrupted his teacher at school (f_6), asks Jack whether he can also have ice cream, presenting him with the situation $X_2 = \{f_2, f_6\}$. As we have seen, H-constraint requires Jack to grant Max's request. Why? Because making his bed means that Max behaved at home (q), while interrupting the teacher means that he misbehaved at school (r'). But then, from the hierarchical standpoint, the situation presented by Max is exactly the same as that presented by Emma: a child who has behaved at home but misbehaved at school. And since Jo has already ruled in favor of ice cream in that situation—on the grounds that behaving at home is a more important reason in favor of ice cream than misbehaving at school is against—precedent requires the same decision from Jack.

So then, why does F-constraint yield a different result? The problem here is that flattening turns the rich, hierarchical justification provided by Jo into something like: *Although she threw her toys on the floor and was inattentive at school, my decision is that Emma can have ice cream because she folded her clothes*. When the justification is flattened in this way, two important pieces of information are lost. The first is that the situation presented by Emma can be described as one in which a child behaved at home but misbehaved at school. The second is that, according to Jo, behaving at home is more important than misbehaving at school. In the absence of this information, the common structure between the situations presented by Emma and Max is obscured, as is the real reason that Jo decided for Emma. Jack is therefore free to rule against Max—Jo's flattened justification does not tell us anything about whether or not making the bed is more important than interrupting one's teacher at school.

[10]The hierarchy $\mathscr{H}$ has the odd feature that the intermediate factors q and q' mirror, respectively, p and p'. Although we could make the hierarchy more realistic by adding more intermediate and base-level factors supporting the application of q and q', we decided to keep the hierarchy $\mathscr{H}$ as it is to avoid introducing unnecessary complications.

Of course, we agree with Bench-Capon that, where authorities—such as Jack and Jo—fail to agree on the structure of the factor hierarchies underlying their decisions, or fail to acknowledge them, then priorities concerning intermediate factors from those hierarchies have little force. But the same holds true of base-level factors. Unless courts agree, at least broadly, on the set of base-level factors bearing on a decision, then previous decisions concerning priorities among these base-level factors are irrelevant. Bench-Capon suggests that there is general agreement on underlying sets of base-level factors: they correspond to "commonly used phrases" or "commonly considered" aspects of a case. We want to suggest that there is, at least very often, broad agreement on the nature of and relations among intermediate factors bearing on higher-level issues—there is at least broad agreement that, in the domain of family life, for example, good behavior at home favors a reward of some kind, that tidying up one's room favors good behavior at home, and that making one's bed favors tidying up one's room.

As noted, Bench-Capon also points out that intermediate factors were originally introduced in CATO only in order to "group together [base-level] factors that could be considered to substitute for or cancel one another, which was a particular requirement of CATO's motivation." We agree that this is why intermediate factors were initially introduced into the representation. But this does not mean that intermediate factors are not real—we should not confuse motivation with metaphysics. The intermediate factor *tidying up one's room*, for example, is not just an empty cipher indicating that the base-level factors *folding clothes* and *making one's bed* can be substituted for each other in certain arguments. Instead, the intermediate factor plays an explanatory role: the reason that *folding clothes* and *making one's bed* are intersubstitutable is that they both favor *tidying up one's room*.

3.2. F-constraint without H-constraint

Cases that are F-constrained but not H-constrained are more problematic. We will not offer a general theory, but simply point out some issues emerging from Bench-Capon's analysis of our example from Figure 2.

The example concerns the case c_3, top left, in which the court decided $X_3 = \{f_1, f_4, f_5\}$ for π on the basis of the hierarchical justification j_3. When a later court is confronted with the new fact situation $X_4 = \{f_1, f_5\}$, the background case base $\Gamma_3 = \{c_3\}$ does not require a decision for π by H-constraint, but $\mathit{flatten}(\Gamma_3)$, bottom left, does require a decision for π by F-constraint. Which of the two results is correct?

Bench-Capon thinks that the later court should not be constrained, but not because he thinks that H-constraint is correct. Instead, he argues that $\mathit{flatten}(\Gamma_3)$ should not require a decision for π in X_4 by F-constraint. The reason is that, given how c_3 was decided, he thinks the $\mathit{flatten}(c_3)$ decision that δ holds should depend on both f_1 and f_4, rather than f_1 alone: f_1 is required because it provides a reason for π and f_4 is required because, as he says, it "neutralizes" the factor f_5 favoring δ. But, if $\mathit{flatten}(c_3)$ is the decision that δ holds because of f_1 and f_4, then this decision tells us that f_1 and f_4 together are more important than f_5, which is perfectly consistent with the $\mathit{flatten}(c_4)$ decision, bottom right, that f_5 is more important than f_1. Bench-Capon concludes that the example from Figure 2 may be "an artifact of the flattening mechanism in [13]."

These arguments raise a number of issues concerning our notion of flattening, which we realize can be defined in different ways. A first issue is whether f_1 alone is the reason for the $\mathit{flatten}(c_3)$ decision. We think that there are concrete interpretations of c_3 in which $\mathit{flatten}(c_3)$ should depend on f_1 alone. Recall our earlier example of Jack and Jo who, in order to decide whether or not to give their children, Emma and Max, a treat, rely on the concrete interpretation of the hierarchy $\mathscr{H}$ from Figure 3. Suppose that one day Emma, who folded her clothes (f_1), turned in her homework (f_4), but was inattentive in class (f_5), asks Jo for ice cream; this situation can be represented through the previous

$X_3 = \{f_1, f_4, f_5\}$. And the case c_3 can then be taken to represent Jo's decision to let Emma have ice cream, where the earlier justification j_3 from that case can be interpreted as follows: *Emma tidied up her room because she folded her clothes and she behaved at school because she turned in her homework. Since she tidied up her room, Emma behaved at home. Because of this, my decision is that Emma can have ice cream.*

What is the correct way to flatten Jo's hierarchical justification? Our notion of flattening turns it into something like: *Although she was inattentive in class, my decision is that Emma can have ice cream because she folded her clothes*. But, according to Bench-Capon's argument, the flattened justification should rather be: *Although she was inattentive in class, my decision is that Emma can have ice cream because she folded her clothes and turned in her homework*. To decide which flattening is correct, we need to ask: is turning in homework part of the reason why the child can have ice cream?

Given our concrete interpretation of the hierarchy $\mathscr{H}$, the natural answer is No: turning in homework does not itself support a decision in favor of ice cream—it simply prevents misbehavior at school from supporting a decision against ice cream. In the domain of family life, this makes perfect sense: parents often expect their children to turn in their homework and they would not give the children a treat simply because they did what they were expected to do—although they may deny a treat if the children did not do what they were expected to do. Using terminology that goes back to Dancy [20], turning in homework can be classified as an *enabler*, rather than a reason: it is an external consideration that allows the reason that a child folded their clothes to support a decision in favor of ice cream, despite not being itself a reason for ice cream. Hence, our notion of flattening—which returns reasons alone, rather than enablers as well—seems to yield the correct result in this case.

But now a second issue arises. Suppose that, after Jo made her decision, Max, who folded his clothes (f_1) but was inattentive in class (f_5), asks Jack whether he can also have ice cream, presenting the situation $X_4 = \{f_1, f_5\}$. From the hierarchical standpoint, nothing prevents Jack from ruling against ice cream because, unlike Emma, Max misbehaved at school. But F-constraint requires Jack to grant Max's request. Why? Because, as we saw above, folding clothes is the only reason justifying Jo's flattened decision—turning in homework is an enabler. According to the standard reason model, folding clothes is thus more important than being inattentive in class, no matter whether homework was turned in or not. The problem is that, since the standard reason model does not contain a distinction between reasons and enablers, it cannot capture the information that, absent its enabler, folding clothes may fail to support ice cream.

One way around this problem might be to refine the standard reason model by introducing a distinction between reasons and enablers that would prevent cases like the one presented by Max from being F-constrained. This may be an interesting way to develop Bench-Capon's intuition that, in the example from Figure 2, the court should not be required to decide X_4 for π by F-constraint. But, at least in our concrete interpretation of the example, the distinction between reasons and enablers seems to be grounded in the factor hierarchy—it is because of the role they play in Jo's hierarchical justification that folding clothes and turning in homework are, respectively, a reason and an enabler. This suggests that, at least in our example, the factor hierarchy has an important role to play. Two natural questions for future work would be, first, whether these considerations can be generalized and, if so, whether enablers can be defined precisely in our setting.

4. Conclusion

Working over many years and with many collaborators, Bench-Capon has developed a clear, detailed, and fruitful vision of the most helpful knowledge representation formalisms for use in AI and Law. One consequence of this vision is the conclusion that,

at least in a purely legal setting, precedential constraint is best defined in terms of hierarchies without intermediate factors. We do not challenge this conclusion. We do not have the legal expertise to do so—although we note again that some researchers with the relevant expertise, like Branting or Prakken and Sartor, have hinted at a contrary position.

Instead, our goal in this paper is to suggest, first of all, that many of the ideas and tools developed in AI and Law—many, indeed, developed by Bench-Capon himself—have important applications outside the purely legal domain, in areas such as Ethics and AI and Explainable AI. And second, that in at least some of these other areas, even if not in the law proper, intermediate factors may help us understand crucial patterns of argument and explanation, and in particular, that decisions concerning applicability of intermediate factors may have precedential import.

References

[1] Canavotto I, Horty J. Reasoning with hierarchies of open-textured predicates. In: Proceedings of the Nineteenth International Conference on AI and Law (ICAIL-23). The Association for Computing Machinery Press; 2023. p. 52–61.
[2] Horty J, Bench-Capon T. A factor-based definition of precedential constraint. Artificial Intelligence and Law. 2012;20:181–214.
[3] Goodhart A. Determining the *ratio decidendi* of a case. Yale Law Journal. 1930;40:161–183.
[4] Laird J, Rosenbloom P, Newell A. Chunking in Soar: the anatomy of a general learning mechanism. Machine Learning. 1986;1:11–46.
[5] Erol K, Hendler J, Nau D. HTN planning: complexity and expressivity. In: Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94). The Association for Computing Machinery Press; 1994. p. 1123–1128.
[6] Bench-Capon T, Atkinson K. Precedential constraint: the role of issues. In: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL 2021). The Association for Computing Machinery Press; 2021. p. 12–21.
[7] Bench-Capon T. A note on hierarchical constraint; 2023. Jurix 2023, submission 10.
[8] Aleven V. Teaching Case-Based Argumentation Through a Model and Examples. Intelligent Systems Program, University of Pittsburgh; 1997.
[9] Brüninghaus S, Ashley K. Predicting outcomes of case based legal arguments. In: Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAIL-03). The Association for Computing Machinery Press; 2003. p. 233–242.
[10] Grabmair M, Ashley K. Faciliting case comparison using value judgments and intermediate legal concepts. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Law (ICAIL-11). The Association for Computing Machinery Press; 2011. p. 161–170.
[11] Branting LK. Reasoning with portions of precedents. In: Proceedings of the Third International Conference on Artificial Intelligence and Law (ICAIL-91). ACM Press; 1991. p. 145–154.
[12] Prakken H, Sartor G. Modelling reasoning with precedents in a formal dialogue game. Artificial Intelligence and Law. 1998;6:231–287.
[13] van Woerkom W, Grossi D, Prakken H, Verheij B. Hierarchical precedential constraint. In: Proceedings of the Nineteenth International Conference on AI and Law (ICAIL-23). The Association for Computing Machinery Press; 2023. p. 333–342.
[14] Canavotto I, Horty J. Piecemeal Knowledge Acquisition for Computational Normative Reasoning. In: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES'2022). ACM Press; 2022. p. 171–180.
[15] Liu X, Lorini E. A unified logical framework for explanations in classifier systems. Journal of Logic and Computation. 2023;33(2):485–515.
[16] Peters J, Bex FJ, Prakken H. Model- and data-agnostic justifications with *a fortiori* case-based argumentation. In: Proceedings of the Nineteenth Internation Conference on Artificial Intelligence and Law (ICAIL-23). The Association for Computing Machinery Press; 2023. p. 207–216.
[17] Prakken H, Ratsma R. A Top-Level Model of Case-Based Argumentation for Explanation: Formalisation and Experiments. Argument & Computation. 2022;13(2):159–194.
[18] Levin H. The food stays in the kitchen: everything I needed to learn about statutory interpretation I learned by the time I was nine. Green Bag. 2009;12:337–344.
[19] Twining W, Miers D. How to Do Things with Rules. Cambridge University Press; 2010. Fifth edition.
[20] Dancy J. Ethics Without Principles. Oxford University Press; 2004.